

**BIOS and Kernel
Developer's Guide
(BKDG)
for AMD Family 15h
Models 10h-1Fh
Processors**

© 2011–2015 Advanced Micro Devices, Inc. All rights reserved.

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

Trademarks

AMD, the AMD Arrow logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

ARM, AMBA, Cortex and TrustZone are registered trademarks of ARM Limited.

HDMI is a trademark of HDMI Licensing, LLC.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

Microsoft and DirectX are registered trademarks of Microsoft Corporation.

MMX is a trademark of Intel Corporation.

PCI Express and PCIe are registered trademarks of PCI-Special Interest Group (PCI-SIG).

Dolby Laboratories, Inc.

Manufactured under license from Dolby Laboratories.

Rovi Corporation

This device is protected by U.S. patents and other intellectual property rights. The use of Rovi Corporation's copy protection technology in the device must be authorized by Rovi Corporation and is intended for home and other limited pay-per-view uses only, unless otherwise authorized in writing by Rovi Corporation.

Reverse engineering or disassembly is prohibited.

USE OF THIS PRODUCT IN ANY MANNER THAT COMPLIES WITH THE MPEG ACTUAL OR DE FACTO VIDEO AND/OR AUDIO STANDARDS IS EXPRESSLY PROHIBITED WITHOUT ALL NECESSARY LICENSES UNDER APPLICABLE PATENTS. SUCH LICENSES MAY BE ACQUIRED FROM VARIOUS THIRD PARTIES INCLUDING, BUT NOT LIMITED TO, IN THE MPEG PATENT PORTFOLIO, WHICH LICENSE IS AVAILABLE FROM MPEG LA, L.L.C., 6312 S. FIDDLERS GREEN CIRCLE, SUITE 400E, GREENWOOD VILLAGE, COLORADO 80111.

Table of Contents

1	Overview	18
1.1	Intended Audience	18
1.2	Reference Documents	18
1.3	Conventions	18
1.3.1	Numbering	18
1.3.2	Arithmetic And Logical Operators	19
1.4	Definitions	20
1.5	Changes Between Revisions and Product Variations	25
1.5.1	Revision Conventions	25
1.5.2	Major Changes Relative to Family 15h Models 00h-0Fh Processors	25
1.5.2.1	Major Changes to Core/NB Performance Counters	26
1.5.3	Changes For Revision RL-A1	26
2	Functional Description	27
2.1	Processor Overview	27
2.2	System Overview	27
2.3	Processor Initialization	28
2.3.1	BSC initialization	28
2.3.2	AP initialization	28
2.3.3	Using L2 Cache as General Storage During Boot	29
2.4	Core	31
2.4.1	Compute Unit	31
2.4.1.1	Registers Shared by Cores in a Compute Unit	31
2.4.2	Virtual Address Space	31
2.4.3	Processor Cores and Downcoring	31
2.4.3.1	Software Downcoring using D18F3x190[DisCore]	32
2.4.4	Physical Address Space	32
2.4.5	System Address Map	32
2.4.5.1	Memory Access to the Physical Address Space	32
2.4.5.1.1	Determining Memory Type	32
2.4.5.1.2	Determining The Access Destination for Core Accesses	33
2.4.6	Timers	33
2.4.7	Implicit Conditions for TLB Invalidation	33
2.4.8	Interrupts	34
2.4.8.1	Local APIC	34
2.4.8.1.1	Detecting and Enabling	34
2.4.8.1.2	APIC Register Space	34
2.4.8.1.3	ApicId Enumeration Requirements	34
2.4.8.1.4	Physical Destination Mode	35
2.4.8.1.5	Logical Destination Mode	35
2.4.8.1.6	Interrupt Delivery	35
2.4.8.1.7	Vectored Interrupt Handling	36
2.4.8.1.8	Interrupt Masking	36
2.4.8.1.9	Spurious Interrupts	36
2.4.8.1.10	Spurious Interrupts Caused by Timer Tick Interrupt	36
2.4.8.1.11	Lowest-Priority Interrupt Arbitration	37
2.4.8.1.12	Inter-Processor Interrupts	37
2.4.8.1.13	APIC Timer Operation	37

2.4.8.1.14	Generalized Local Vector Table	37
2.4.8.1.15	State at Reset	38
2.4.8.2	System Management Mode (SMM)	38
2.4.8.2.1	SMM Overview	38
2.4.8.2.2	Operating Mode and Default Register Values	38
2.4.8.2.3	SMI Sources And Delivery	39
2.4.8.2.4	SMM Initial State	39
2.4.8.2.5	SMM Save State	40
2.4.8.2.6	Exceptions and Interrupts in SMM	45
2.4.8.2.7	The Protected ASeg and TSeg Areas	45
2.4.8.2.8	SMM Special Cycles	45
2.4.8.2.9	Locking SMM	46
2.4.8.2.10	Synchronizing SMM Entry (Spring-Boarding)	46
2.4.9	Secure Virtual Machine Mode (SVM)	47
2.4.9.1	BIOS support for SVM Disable	47
2.4.10	CPUID Instruction	48
2.4.10.1	Multi-Core Support	48
2.5	Power Management	49
2.5.1	Processor Power Planes And Voltage Control	49
2.5.1.1	Serial VID Interface	49
2.5.1.1.1	SVI2 Features	49
2.5.1.2	Internal VID Registers and Encodings	50
2.5.1.2.1	MinVid and MaxVid Check	50
2.5.1.3	Low Power Features	50
2.5.1.3.1	PSIx_L Bit	50
2.5.1.3.1.1	BIOS Requirements for PSIO_L	50
2.5.1.3.1.2	BIOS Requirements for PSII_L	52
2.5.1.3.2	Low Power Voltages	52
2.5.1.4	Voltage Transitions	52
2.5.1.4.1	Hardware-Initiated Voltage Transitions	52
2.5.1.4.2	Software-Initiated Voltage Transitions	52
2.5.1.4.2.1	Software-Initiated NB Voltage Transitions	52
2.5.1.4.2.2	Software-Initiated Core Voltage Transitions	53
2.5.2	Frequency and Voltage Domain Dependencies	53
2.5.2.1	Dependencies Between Cores	53
2.5.2.2	Dependencies Between Subcomponents on VDDNB	54
2.5.2.3	BIOS Requirements for Power Plane Initialization	54
2.5.3	CPU Power Management	54
2.5.3.1	Core P-states	54
2.5.3.1.1	Application Power Management (APM)	54
2.5.3.1.2	Core P-state Naming and Numbering	55
2.5.3.1.2.1	Software P-state Numbering	55
2.5.3.1.2.2	Hardware P-state Numbering	56
2.5.3.1.3	Core P-state Control	56
2.5.3.1.4	Core P-state Visibility	57
2.5.3.1.5	Core P-state Limits	57
2.5.3.1.6	Core P-state Transition Behavior	57
2.5.3.1.7	BIOS Requirements for Core P-state Initialization and Transitions	58
2.5.3.1.8	Processor-Systemboard Power Delivery Compatibility Check	59
2.5.3.1.9	BIOS COF and VID Requirements After Warm Reset	60
2.5.3.1.9.1	Core Maximum P-state Transition Sequence After Warm Reset	61

2.5.3.1.9.2	Core Minimum P-state Transition Sequence After Warm Reset	61
2.5.3.1.9.3	ACPI Processor P-state Objects	61
2.5.3.1.9.4	Fixed ACPI Description Table (FADT) Entries	63
2.5.3.1.9.5	XPSS (Microsoft® Extended PSS) Object	63
2.5.3.2	Core C-states	63
2.5.3.2.1	C-state Names and Numbers	63
2.5.3.2.2	C-state Request Interface	63
2.5.3.2.3	C-state Actions	64
2.5.3.2.3.1	C-state Probes and Cache Flushing	64
2.5.3.2.3.2	Core C1 (CC1) State	64
2.5.3.2.3.3	Core C6 (CC6) State	64
2.5.3.2.3.4	Package C6 (PC6) State	65
2.5.3.2.4	C-state Request Monitors	65
2.5.3.2.4.1	FCH Messaging	65
2.5.3.2.4.2	Timer Tick Monitor	66
2.5.3.2.4.3	Cache Flush On Halt Saturation Counter	66
2.5.3.2.5	Exiting C-states	66
2.5.3.2.6	ACPI Processor C-state Objects	66
2.5.3.2.6.1	_CST	66
2.5.3.2.6.2	_CSD	67
2.5.3.2.6.3	_CRS	67
2.5.3.2.6.4	Fixed ACPI Description Table (FADT) Entries	67
2.5.3.2.7	BIOS Requirements for Initialization	67
2.5.3.3	Effective Frequency	67
2.5.4	NB Power Management	68
2.5.4.1	NB P-states	68
2.5.4.1.1	NB P-state Transitions	68
2.5.4.1.2	BIOS NB P-state Configuration	69
2.5.4.1.2.1	NB P-state COF and VID Synchronization After Warm Reset	69
2.5.4.1.2.2	NB P-state Voltage Adjustment for MEMCLK	70
2.5.4.1.2.3	NB P-state Configuration for Runtime	70
2.5.4.2	NB C-states	70
2.5.5	P-state Bandwidth Requirements	71
2.5.6	GPU and Root Complex Power Management	71
2.5.6.1	Dynamic Power Management (DPM)	71
2.5.6.1.1	Activity Monitors	72
2.5.6.1.2	SCLK DPM	72
2.5.6.1.3	LCLK DPM	72
2.5.6.2	GPU and Root Complex Power Gating	72
2.5.7	DRAM Power Management	73
2.5.7.1	Memory P-states	73
2.5.7.2	DRAM Self-Refresh	73
2.5.7.3	Stutter Mode	74
2.5.7.3.1	System BIOS Requirements for Stutter Mode Operation During POST	74
2.5.7.4	EVENT_L	74
2.5.8	System Power Management	74
2.5.8.1	S-states	75
2.5.8.1.1	ACPI Suspend to RAM State (S3)	75
2.5.9	Bidirectional Application Power Management (BAPM)	75
2.5.9.1	Hybrid Boost	75
2.6	Performance Monitoring	77

2.6.1	Core Performance Monitor Counters	77
2.6.2	NB Performance Monitor Counters	78
2.6.3	Instruction Based Sampling (IBS)	79
2.7	Configuration Space	80
2.7.1	MMIO Configuration Coding Requirements	80
2.7.2	MMIO Configuration Ordering	80
2.7.3	Processor Configuration Space	81
2.8	Northbridge (NB)	82
2.8.1	NB Architecture	82
2.8.2	NB Routing	82
2.8.2.1	Address Space Routing	82
2.8.2.1.1	DRAM and MMIO Memory Space	82
2.8.2.1.2	IO Space	83
2.8.2.1.3	Configuration Space	83
2.9	DRAM Controllers (DCTs)	84
2.9.1	DCT Configuration Registers	87
2.9.2	DDR Pad to Processor Pin Mapping	87
2.9.2.1	DDR Chip to Pad Mapping	88
2.9.3	DRAM Controller Direct Response Mode	89
2.9.4	DRAM Data Burst Mapping	89
2.9.5	DCT/DRAM Initialization and Resume	90
2.9.5.1	Low Voltage DDR3	90
2.9.5.2	NB P-state Specific Configuration	91
2.9.5.3	Memory P-state Specific Configuration	92
2.9.5.4	DDR Phy Initialization	92
2.9.5.4.1	Phy Voltage Level Programming	93
2.9.5.4.2	DRAM Channel Frequency Change	93
2.9.5.4.2.1	Requirements for DRAM Frequency Change During Training	94
2.9.5.4.3	Phy Fence Programming	94
2.9.5.4.3.1	Phy Fence Training	95
2.9.5.4.4	Phy Compensation Initialization	96
2.9.5.5	SPD ROM-Based Configuration	101
2.9.5.6	Non-SPD ROM-Based Configuration	102
2.9.5.6.1	TrdrdSdSc, TrdrdSdDd, and TrdrdDd (Read to Read Timing)	103
2.9.5.6.2	TwrwrSdSc, TwrwrSdDc, TwrwrDd (Write to Write Timing)	103
2.9.5.6.3	Twrrd (Write to Read DIMM Termination Turn-around)	104
2.9.5.6.4	TrwtTO (Read-to-Write Turnaround for Data, DQS Contention)	104
2.9.5.6.5	DRAM ODT Control	104
2.9.5.6.6	DRAM Address Timing and Output Driver Compensation Control	105
2.9.5.7	DCT Training Specific Configuration	114
2.9.5.8	DRAM Device and Controller Initialization	115
2.9.5.8.1	Software DDR3 Device Initialization	115
2.9.5.8.1.1	DDR3 MR Initialization	116
2.9.5.9	DRAM Training	118
2.9.5.9.1	Write Levelization Training	118
2.9.5.9.1.1	Write Leveling Seed Value	120
2.9.5.9.2	DQS Receiver Enable Training	120
2.9.5.9.2.1	DQS Receiver Enable Training Seed Value	121
2.9.5.9.3	DQS Receiver Enable Cycle Training	122
2.9.5.9.4	DQS Position Training	122
2.9.5.9.5	Calculating MaxRdLatency	125

2.9.5.9.5.1	MaxRdLatency Training	126
2.9.5.9.6	Continuous Pattern Generation	126
2.9.5.9.6.1	DRAM Training Pattern Generation	126
2.9.5.10	DRAM Channel Disable	129
2.9.5.11	DRAM Phy Power Savings	129
2.9.6	Memory Interleaving Modes	130
2.9.6.1	Chip Select Interleaving	130
2.9.6.2	Channel Interleaving	132
2.9.7	Memory Hoisting	132
2.9.7.1	DramHoleOffset Programming	132
2.9.7.2	DctSelBaseOffset Programming	133
2.9.8	DRAM CC6/PC6 Storage	134
2.9.9	DRAM On DIMM Thermal Management and Power Capping	134
2.10	Thermal Functions	136
2.10.1	The Tctl Temperature Scale	136
2.10.2	Temperature Slew Rate Control	137
2.10.3	Sideband Temperature Sensor Interface (SB-TSI)	137
2.10.4	Temperature-Driven Logic	137
2.10.4.1	PROCHOT_L and Hardware Thermal Control (HTC)	137
2.10.4.2	Local Hardware Thermal Control (LHTC)	138
2.10.4.3	Software P-state Limit Control	138
2.10.4.4	THERMTRIP	138
2.11	Root Complex	139
2.11.1	Overview	139
2.11.2	Interrupt Routing	139
2.11.3	Links	140
2.11.3.1	Overview	140
2.11.3.2	Link Configurations	140
2.11.3.3	Clocking	143
2.11.4	Root Complex Configuration	143
2.11.4.1	LPC MMIO Requirements	143
2.11.4.2	Link Configuration and Initialization	144
2.11.4.2.1	Clock Configuration	144
2.11.4.2.2	Link Configuration and Core Initialization	144
2.11.4.2.3	Link Training	145
2.11.4.3	Miscellaneous Features	145
2.11.4.3.1	Straps	145
2.11.4.3.2	Lane Reversal	145
2.11.4.3.3	Link Speed Changes	146
2.11.4.3.3.1	Software Initiated Link Speed Changes	146
2.11.4.3.3.2	Autonomous Link Speed Changes	146
2.11.4.3.4	Deemphasis	146
2.11.4.4	Power Management	146
2.11.4.4.1	Link States	146
2.11.4.4.2	Dynamic Link-width Control	147
2.11.4.5	Link Test and Debug Features	147
2.11.4.5.1	Compliance Mode	147
2.11.5	BIOS Timer	147
2.12	IOMMU	147
2.12.1	IOMMU Configuration Space	147
2.12.2	IOMMU Initialization	148

2.13	System Management Unit (SMU)	149
2.13.1	Software Interrupts	149
2.14	Graphics Processor (GPU)	150
2.14.1	GPU PCI Interface	150
2.14.2	Graphics Memory Controller (GMC)	150
2.14.3	Frame Buffer (FB)	150
2.15	RAS Features	151
2.15.1	Machine Check Architecture	151
2.15.1.1	Machine Check Registers	151
2.15.1.2	Machine Check Errors Classes	152
2.15.1.3	Error Detection, Action, Logging, and Reporting	153
2.15.1.3.1	MCA conditions that cause Shutdown	154
2.15.1.3.2	Error Logging During Overflow	154
2.15.1.4	MCA Initialization	155
2.15.1.5	Error Code	155
2.15.1.6	Handling Machine Check Exceptions	157
2.15.1.6.1	MCA Differentiation Between System-Fatal and Process-Fatal Errors	159
2.15.1.7	Error Thresholding	159
2.15.1.8	Error Diagnosis	160
2.15.1.8.1	Common Diagnosis Information	161
2.15.2	Error Injection and Simulation	162
3	Registers	163
3.1	Register Descriptions and Mnemonics	163
3.1.1	Northbridge MSRs In Multi-Core Products	165
3.1.2	Software Recommendation (BIOS, SBIOS, CBIOS, etc.)	166
3.1.3	Mapping Tables	166
3.1.3.1	Register Mapping	166
3.1.3.2	Index Mapping	166
3.1.3.3	Field Mapping	166
3.1.3.4	Broadcast Mapping	166
3.1.3.5	Valid Values	167
3.2	IO Space Registers	167
3.3	Device 0 Function 0 (Root Complex) Configuration Registers	168
3.4	Device 0 Function 2 (IOMMU) Configuration Registers	236
3.5	Device 1 Function 0 Configuration Registers	268
3.6	Device 1 Function 1 (Audio Controller) Configuration Registers	280
3.7	Device [8:2] Function 0 (Root Port) Configuration Registers	291
3.8	Device 18h Function 0 Configuration Registers	320
3.9	Device 18h Function 1 Configuration Registers	328
3.10	Device 18h Function 2 Configuration Registers	338
3.11	Device 18h Function 3 Configuration Registers	419
3.12	Device 18h Function 4 Configuration Registers	446
3.13	Device 18h Function 5 Configuration Registers	453
3.14	GPU Memory Mapped Registers	464
3.15	IOMMU Memory Mapped Registers	467
3.16	APIC Registers	482
3.17	CPUID Instruction Registers	493
3.18	MSRs - MSR0000_0xxx	519
3.19	MSRs - MSRC000_0xxx	562
3.20	MSRs - MSRC001_0xxx	567

3.21	MSRs - MSRC001_1xxx	592
3.22	Core Performance Counter Events	606
3.22.1	PMCx0[1F:00] Events (FP)	606
3.22.2	PMCx0[3F:20] Events (LS)	607
3.22.3	PMCx0[5F:40] Events (DC)	609
3.22.4	PMCx[1,0][7F:60] Events (CU)	612
3.22.5	PMCx[1,0][9F:80] Events (IC)	615
3.22.6	PMCx0[BF:A0] Events	617
3.22.7	PMCx[1,0][DF:C0] Events (EX, DE)	617
3.23	NB Performance Counter Events	622
3.23.1	0E[7:0] Events (Memory Controller)	622
3.23.2	0E[F:8] Events (Crossbar)	624
3.23.3	0F[F:0] Events (Crossbar)	626
3.23.4	1E[F:0] Events (Crossbar)	626
3.23.5	1F[F:0] Events (Memory Controller, Crossbar)	630
4	Register List	631

List of Figures

Figure 1:	A processor	27
Figure 2:	System Diagram.....	28
Figure 3:	DQS Position Training Example Results	124
Figure 4:	DQS Position Training Insertion Delay Recovery Example Results.....	125
Figure 5:	Example Cases for Programming DramHoleOffset.....	133
Figure 6:	Example Cases for Programming DctSelBaseOffset.....	134
Figure 7:	Tctl scale	136
Figure 8:	Root complex topology	139
Figure 9:	Phy clock configuration	143
Figure 10:	Phy recovered clock and sample clock	213
Figure 11:	Address/Command Timing at the Processor Pins.....	353

List of Tables

Table 1:	Arithmetic and Logical Operators.....	19
Table 2:	Functions.....	19
Table 3:	Definitions.....	20
Table 4:	Processor revision conventions.....	25
Table 5:	SMM Initial State.....	39
Table 6:	SMM Save State.....	40
Table 7:	Power Management Support.....	49
Table 8:	Software P-state Naming.....	55
Table 9:	Software P-state Control.....	56
Table 10:	Core PMC mapping to PERF_CTL[5:0].....	77
Table 11:	DCT Definitions.....	84
Table 12:	DDR3 unbuffered DIMM maximum frequency support for FM2 (per channel).....	85
Table 13:	DDR3 SO-DIMM maximum frequency support for FS1r2 (per channel).....	85
Table 14:	DDR3 unbuffered DIMM maximum frequency support for FP2 (per channel) for (TN-A0 TN-A1)	85
Table 15:	DDR3 SO-DIMM maximum frequency support for FP2 (per channel) for (TN-A0 TN-A1).	86
Table 16:	DDR3 unbuffered DIMM maximum frequency support for FP2 (per channel) for RL-A1.....	86
Table 17:	DDR3 SO-DIMM maximum frequency support for FP2 (per channel) for RL-A1.....	86
Table 18:	DDR3 Soldered-down DRAM maximum frequency support for FP2 (per channel).....	86
Table 19:	Package pin mapping.....	87
Table 20:	Pad (from chiplet) pin mapping.....	89
Table 21:	DDR PLL Lock Time.....	94
Table 22:	Phy predriver calibration codes for Data/DQS at 1.5V.....	96
Table 23:	Phy predriver calibration codes for Data/DQS at 1.35V.....	97
Table 24:	Phy Predriver Calibration Codes for Data/DQS at 1.25V.....	97
Table 25:	Phy predriver calibration codes for Cmd/Addr at 1.5V.....	98
Table 26:	Phy predriver calibration codes for Cmd/Addr at 1.35V.....	99
Table 27:	Phy Predriver Calibration Codes for Cmd/Addr at 1.25V.....	99
Table 28:	Phy predriver calibration codes for Clock at 1.5V.....	100
Table 29:	Phy predriver calibration codes for Clock at 1.35V.....	100
Table 30:	Phy Predriver Calibration Codes for Clock at 1.25V.....	101
Table 31:	DDR3 DIMM ODT Pattern.....	105
Table 32:	BIOS recommendations for UDIMM address timings and output driver control for FM2 or FS1r2 Package.....	105
Table 33:	BIOS recommendations for UDIMM address timings and output driver control for FP2 Package.....	108
Table 34:	BIOS recommendations for SO-DIMM address timings and output driver control for FM2 or FS1r2 Package.....	109
Table 35:	BIOS recommendations for SO-DIMM address timings and output driver control for FP2 Package.....	112
Table 36:	BIOS recommendations for DRAM soldered down address timings and output driver control for FP2 Package.....	113

Table 37:	DCT Training Specific Register Values	114
Table 38.	DDR3 MR0	116
Table 39.	DDR3 MR1	117
Table 40.	DDR3 MR2	117
Table 41.	DDR3 MR3	118
Table 42.	DDR3 Write Leveling Seed Values.....	120
Table 43.	DDR3 DQS Receiver Enable Training Seed Values.....	122
Table 44.	Command Generation and Data Comparison	129
Table 45.	DDR3 Swapped Normalized Address Lines for CS Interleaving.....	130
Table 46.	Example storage region configuration	134
Table 47:	INTx Mapping.....	139
Table 48:	Lane Id Mapping.....	140
Table 49:	Supported Gfx Port Configurations	140
Table 50:	Supported Gfx Port Configurations for D0F0xE4_x013[3:1]_804[3:0].....	141
Table 51:	Supported DP0/DP1 DDI Link Configurations	142
Table 52:	Supported DP2 DDI Link Configurations	142
Table 53:	Supported General Purpose (GPP) Link Configurations	142
Table 54:	BIOS Services.....	149
Table 55:	Recommended Frame Buffer Configurations	150
Table 56:	MCA Register Cross-Reference Table.....	152
Table 57:	Overwrite Priorities for All Banks	155
Table 58:	Error Code Types	156
Table 59:	Error Codes: Transaction Type	156
Table 60:	Error codes: cache level	156
Table 61:	Error Codes: Memory Transaction Type.....	156
Table 62:	Error Codes: Participation Processor.....	157
Table 63:	Error Codes: Memory or IO.....	157
Table 64:	Error Scope Hierarchy	159
Table 65:	Registers Commonly Used for Diagnosis.....	161
Table 66:	Terminology in Register Descriptions	164
Table 67:	BIOS recommendations for D0F0x64_x5[B,9,7,5,3,1]	176
Table 68:	Register Mapping for D0F0xBC_x1F2[E0:00:step20].....	182
Table 69:	Register Mapping for D0F0xBC_x1F2[E8:08:step20].....	182
Table 70:	Register Mapping for D0F0xBC_x1F2[F0:10:step20].....	183
Table 71:	Index addresses for D0F0xE4_x0[2:1]1[3:0]_0010	201
Table 72:	Index addresses for D0F0xE4_x0[2:1]1[3:0]_0011	202
Table 73:	Index addresses for D0F0xE4_x0[2:1]1[3:0]_001[3:2].....	203
Table 74:	Index addresses for D0F0xE4_x0[2:1]1[3:0]_0015	204
Table 75:	Per phy register addresses to pin mappings	205
Table 76:	Per nibble register addresses to pin mappings	205
Table 77:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_0000	205
Table 78:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000[2:1].....	206
Table 79:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000[A:9].....	207
Table 80:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000[C:B]	208

Table 81:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000D	209
Table 82:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_2000	209
Table 83:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_2002	210
Table 84:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_2005	210
Table 85:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_2008	211
Table 86:	Phy per receiver lane register addresses	212
Table 87:	Phy receiver broadcast register addresses	212
Table 88:	Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]1	213
Table 89:	Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]1	213
Table 90:	Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]2	214
Table 91:	Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]2	214
Table 92:	Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]5	216
Table 93:	Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]5	216
Table 94:	Recommended DCV settings	216
Table 95:	Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]6	217
Table 96:	Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]6	217
Table 97:	Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]A	217
Table 98:	Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]A	218
Table 99:	Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]F	219
Table 100:	Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]F	219
Table 101:	Phy per transmitter lane register addresses	220
Table 102:	Phy transmitter broadcast register addresses	220
Table 103:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]0	220
Table 104:	Broadcast Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]0	221
Table 105:	Recommended link configuration	221
Table 106:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5	221
Table 107:	Broadcast Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5	222
Table 108:	Recommended preemphasis settings	222
Table 109:	Index Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]6	223
Table 110:	Broadcast Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]6	223
Table 111:	Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[D:C][7:0][8,0]5	223
Table 112:	Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[D:C][7:0][8,0]5	223
Table 113:	Mapping for wrapper registers	224
Table 114:	Index address mapping for D0F0xE4_x013[1:0]_0[C:8]00	225
Table 115:	Index address mapping for D0F0xE4_x013[1:0]_0[C:8]03	225
Table 116:	Reserved field mappings for D0F0xE4_x013[3:0]_8013	226
Table 117:	Reserved field mappings for D0F0xE4_x013[3:0]_8014	227
Table 118:	Reserved field mappings for D0F0xE4_x013[3:0]_8016	229
Table 119:	Index address mapping for D0F0xE4_x013[3:1]_804[3:0]	234
Table 120:	Register Stream mappings for D0F0xE4_x013[3:1]_804[E:8]	234
Table 121:	Valid values for D0F2xFC_x20_L1sel[3:0]	264
Table 122:	Reset mapping for D[8:2]F0x00	291
Table 123:	Link controller state encodings	311
Table 124:	Register Mapping for D18F1x[1B8,1B0,1A8,1A0,B8,B0,A8,A0,98,90,88,80]	330

Table 125:	Register Mapping for D18F1x[1BC,1B4,1AC,1A4,BC,B4,AC,A4,9C,94,8C,84]	331
Table 126:	Register Mapping for D18F1x[D18F1x[1CC:1C0,19C:180]	332
Table 127:	Register Mapping for D18F1x[D8,D0,C8,C0]	333
Table 128:	Register Mapping for D18F1x[DC,D4,CC,C4]	334
Table 129:	Register Mapping for D18F1x[EC:E0].....	335
Table 130:	DIMM, Chip Select, and Register Mapping	338
Table 131:	DDR3 DRAM Address Mapping.....	342
Table 132:	Memory Clock Frequency Value Definition.....	345
Table 133:	Memory P-state Specific Indirect CSR Access by non-BIOS Software.....	349
Table 134:	Memory P-state Specific DCT CSR In Indirect Space	350
Table 135:	Index Mapping for D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0].....	351
Table 136:	Byte Lane Mapping for D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0]	352
Table 137:	Index Mapping for D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0].....	354
Table 138:	Byte Lane Mapping for D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0]	354
Table 139:	Index Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0]	358
Table 140:	Byte Lane Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0].....	359
Table 141:	Index Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0]	359
Table 142:	Byte Lane Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0].....	360
Table 143:	Register Mapping for D18F2x9C_x0000_00[51:50]_dct[1:0]	361
Table 144:	Byte Lane Mapping for D18F2x9C_x0000_00[51:50]_dct[1:0].....	361
Table 145:	Index Mapping for D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0]	362
Table 146:	Index Mapping for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0].....	363
Table 147:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0]	363
Table 148:	Valid Values for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP].....	363
Table 149:	Index Mapping for D18F2x9C_x0D0F_0[F,7:0]04_dct[1:0]_mp[1:0]	364
Table 150:	Index Mapping for D18F2x9C_x0D0F_0[F,7:0]06_dct[1:0]	365
Table 151:	Index Mapping for D18F2x9C_x0D0F_0[F,7:0]0A_dct[1:0]	365
Table 152:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0]	365
Table 153:	Index addresses for D18F2x9C_x0D0F_0[F,7:0]0F_dct[1:0].....	366
Table 154:	Broadcast write index address for D18F2x9C_x0D0F_0[F,7:0]0F_dct[1:0]	366
Table 155:	Index Addresses for D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0]	366
Table 156:	Index Addresses for D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0]	367
Table 157:	Index Addresses for D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0]	368
Table 158:	Index Addresses for D18F2x9C_x0D0F_0[F,7:0]1C_dct[1:0]_mp[1:0]	369
Table 159:	Index Addresses for D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0].....	369
Table 160:	Index Addresses for D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0].....	370
Table 161:	Broadcast Write Index Address for D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0]	370
Table 162:	Index addresses for D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0]	371
Table 163:	Broadcast write index address for D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0]	371
Table 164:	Index Addresses for D18F2x9C_x0D0F_2[F,2:0]00_dct[1:0]_mp[1:0]	372
Table 165:	Index Mapping for D18F2x9C_x0D0F_2[2:0]02_dct[1:0]	372
Table 166:	Index address mapping for D18F2x9C_x0D0F_[C,8,2][2:0]1F.....	373
Table 167:	Index Addresses for D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0].....	373
Table 168:	Index Mapping for D18F2x9C_x0D0F_8[1:0]0[8,4,0]_dct[1:0]_mp[1:0].....	376

Table 169:	Index Mapping for D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0].....	376
Table 170:	Index Mapping for D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0].....	377
Table 171:	Index Mapping for D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0]	377
Table 172:	Index Mapping for D18F2x9C_x0D0F_[C,8][F:0]30_dct[1:0].....	379
Table 173:	Index Mapping for D18F2x9C_x0D0F_C0[10,0C,08,04]_dct[1:0]_mp[1:0]	380
Table 174:	Index Mapping for D18F2x9C_x0D0F_C020_dct[1:0]_mp[1:0]	380
Table 175:	Index Mapping for D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0]	381
Table 176:	D18F2x1B[4:0] Recommended Settings	395
Table 177:	BIOS Recommendations for RdPtrInit	401
Table 178:	XBAR Definitions.....	426
Table 179:	SMAF Action Definition	429
Table 180:	Buffer Count Definitions	443
Table 181:	D18F5x80[Enabled, DualCore] Definition.....	454
Table 182:	Register Mapping for D18F5x1[6C:60].....	458
Table 183:	NB P-state Definitions	458
Table 184:	Block to register mapping for IOMMUx4[1,0][3:0]00.....	477
Table 185:	Block to register mapping for IOMMUx4[1,0][3:0]04.....	478
Table 186:	Block to register mapping for IOMMUx4[1,0][3:0]08.....	478
Table 187:	Block to register mapping for IOMMUx4[1,0][3:0]10.....	478
Table 188:	Block to register mapping for IOMMUx4[1,0][3:0]14.....	479
Table 189:	Block to register mapping for IOMMUx4[1,0][3:0]18.....	479
Table 190:	Block to register mapping for IOMMUx4[1,0][3:0]1C	480
Table 191:	Block to register mapping for IOMMUx4[1,0][3:0]20.....	480
Table 192:	Block to register mapping for IOMMUx4[1,0][3:0]24.....	480
Table 193:	Block to register mapping for IOMMUx4[1,0][3:0]28.....	481
Table 194:	Block to register mapping for IOMMUx4[1,0][3:0]2C	481
Table 195:	Register Mapping for APIC[170:100]	484
Table 196:	Register Mapping for APIC[1F0:180]	485
Table 197:	Register Mapping for APIC[270:200]	485
Table 199:	Register Mapping for APIC3[60:50]	488
Table 200:	Div[3,1:0] Value Table.....	490
Table 201:	Register Mapping for APIC[4F0:480]	491
Table 202:	Register Mapping for APIC[530:500]	491
Table 203:	CPUID Fn0000_0000_E[B,C,D]X Value	493
Table 204:	CPUID Fn8000_0000_E[B,C,D]X Value	500
Table 205:	Value mapping for CPUID Fn8000_000[4:2]_E[D,C,B,A]X Value.....	503
Table 206:	ECX mapping to Cache Type for CPUID Fn8000_001D_E[D,C,B,A]X.....	512
Table 207:	Register Mapping for MSR0000_020[E,C,A,8,6,4,2,0]	523
Table 208:	Memory Type Definition	523
Table 209:	Register Mapping for MSR0000_020[F,D,B,9,7,5,3,1].....	524
Table 210:	Register Mapping for MSR0000_02[6F:68,59:58,50].....	525
Table 211:	Fixed-size MTRR size and Range Mapping	525
Table 212:	LS Error Descriptions	529
Table 213:	LS Error Signatures.....	530

Table 214:	LS Address Register.....	530
Table 215:	IF Error Descriptions	533
Table 216:	IF Error Signatures.....	536
Table 217:	IF Address Register.....	537
Table 218:	MB, SBU, and SBC Definitions	541
Table 219:	CU Error Descriptions	541
Table 220:	CU Error Signatures.....	543
Table 221:	CU Address Register.....	545
Table 222:	NB Error Descriptions	549
Table 223:	NB Error Signatures, Part 1	550
Table 224:	NB Error Signatures, Part 2	551
Table 225:	NB Address Register Default Encoding	552
Table 226:	NB Address Register for Protocol Errors	552
Table 227:	Protocol Error Type.....	552
Table 228:	NB Address Register for NB Array Errors	553
Table 229:	NB Address Register for Watchdog Timer Errors	553
Table 230:	EX Error Descriptions.....	557
Table 231:	EX Error Signatures	558
Table 232:	EX Address Register.....	558
Table 233:	FP Error Descriptions.....	560
Table 234:	FP Error Signatures.....	561
Table 235:	Register Mapping for MSRC001_00[03:00]	567
Table 236:	Register Mapping for MSRC001_00[07:04]	567
Table 237:	Register Mapping for MSRC001_00[35:30]	572
Table 238:	BIOS recommendation for MSRC001_00[35:30]	573
Table 239:	Register Mapping for MSRC001_00[53:50]	577
Table 240:	Register Mapping for MSRC001_00[6B:64].....	580
Table 241:	P-state Definitions.....	580
Table 242:	Register Mapping for MSRC001_020[A,8,6,4,2,0].....	587
Table 243:	Register Mapping for MSRC001_020[B,9,7,5,3,1].....	589
Table 244:	Register Mapping for MSRC001_024[6,4,2,0]	589
Table 245:	Register Mapping for MSRC001_024[7,5,3,1]	591
Table 246:	Register Mapping for PMCx0D[F:C]	620

Revision History

Revision 3.12 Changes

- 2.10.4.2 [Local Hardware Thermal Control (LHTC)]: Clarified.

Revision 3.10 Changes

- 2.5.3.1.1 [Application Power Management (APM)]: Clarified.
- Table 50: [Supported Gfx Port Configurations for D0F0xE4_x013[3:1]_804[3:0]]: Updated.
- 2.9 [DRAM Controllers (DCTs)]: Updated.
- 2.10.1 [The Tctl Temperature Scale]: Updated.
- 2.14.3 [Frame Buffer (FB)]: Updated.
- Table 131: Updated.
- D18F2x404_dct[1:0][UrMctTokenLimit]: Updated.
- D18F3x64 [Hardware Thermal Control (HTC)], D18F3x68 [Software P-state Limit]: Updated.
- MSRC001_1021[DisWayFilter]: Added .
- MSRC001_1030[IbsFetchVal]: Updated.
- D18F2xA8_dct[1:0][BankSwap]: Corrected.
- 1.5.1 [Revision Conventions]: Updated.
- 1.5.3 [Changes For Revision RL-A1]Added.
- 2.5.4.1.1 [NB P-state Transitions]: Updated.
- 2.9.5.6.6 [DRAM Address Timing and Output Driver Compensation Control]: Updated.
- 2.12 [IOMMU]: Updated.
- D18F2x110[DctSelIntLvAddr[1:0]]: Updated.
- D18F2x114[DctSelIntLvAddr[2]]: Updated.
- MSRC001_0071: Updated.
- MSRC001_024[6,4,2,0][IntCoreSel, IntCoreEn]: Added.

Revision 3.00 Changes

- Initial public release.

1 Overview

The AMD Family 15h Models 10h-1Fh processor (in this document referred to as *the processor*) is a processing unit that supports x86-based instruction sets. The processor includes (a) up to two independent processing units referred to as compute units (each compute unit containing two cores), (b) one PCIe[®] root complex (in this document referred to as the root complex or RC) with generation 2 link support, and (c) up to 2 system memory DRAM interfaces.

AMD Family 15h Models 10h-1Fh processors are distinguished by the combined ExtFamily and BaseFamily fields of the CPUID instruction (see [CPUID Fn8000_0001_EAX](#) in 3.17 [CPUID Instruction Registers]).

1.1 Intended Audience

This document provides the processor behavioral definition and associated design notes. It is intended for platform designers and for programmers involved in the development of low-level BIOS (basic input/output system) functions, drivers, and operating system kernel modules. It assumes prior experience in personal computer platform design, microprocessor programming, and legacy x86 and AMD64 microprocessor architecture. The reader should also have familiarity with various platform technologies, such as DDR DRAM.

1.2 Reference Documents

- Advanced Configuration and Power Interface (ACPI) Specification. www.acpi.info.
- AMD64 Architecture Programmer's Manual Volume 1: Application Programming, #24592.
- AMD64 Architecture Programmer's Manual Volume 2: System Programming, #24593.
- AMD64 Architecture Programmer's Manual Volume 3: Instruction-Set Reference, #24594.
- AMD64 Architecture Programmer's Manual Volume 4: 128-Bit and 256-Bit Media Instructions, #26568.
- AMD64 Architecture Programmer's Manual Volume 5: 64-Bit Media and x87 Floating-Point Instructions, #26569.
- AMD I/O Virtualization Technology (IOMMU) Specification, #34434.
- Software Optimization Guide for AMD Family 15h Processors, #47414.
- CPUID Specification, #25481.
- Electrical Data Sheet for AMD Family 15h Models 10h-1Fh Processors, #47080.
- Revision Guide for AMD Family 15h Models 10h-1Fh Processors, #48931.
- JEDEC standards. www.jedec.org.
- PCI local bus specification. (www.pcisig.org).
- PCI Express specification. (www.pcisig.org).
- System Management Bus (SMBus) specification. (www.smbus.org).
- AMD64 Technology Lightweight Profiling Specification, #43724.

1.3 Conventions

1.3.1 Numbering

- **Binary numbers.** Binary numbers are indicated by appending a “b” at the end, e.g., 0110b.
- **Decimal numbers.** Unless specified otherwise, all numbers are decimal. This rule does not apply to the register mnemonics described in 3.1 [Register Descriptions and Mnemonics]; register mnemonics all utilize hexadecimal numbering.
- **Hexadecimal numbers.** hexadecimal numbers are indicated by appending an “h” to the end, e.g., 45f8h.
- **Underscores in numbers.** Underscores are used to break up numbers to make them more readable. They do not imply any operation. E.g., 0110_1100b.

1.3.2 Arithmetic And Logical Operators

In this document, formulas follow some Verilog conventions for logic equations.

Table 1: Arithmetic and Logical Operators

Operator	Definition
{}	Curly brackets are used to indicate a group of bits that are concatenated together. Each set of bits is separated by a comma. E.g., {Addr[3:2], Xlate[3:0]} represents a 6-bit value; the two MSBs are Addr[3:2] and the four LSB's are Xlate[3:0].
	Bitwise OR operator. E.g. (01b 10b == 11b).
	Logical OR operator. E.g. (01b 10b == 01b).
&	Bitwise AND operator. E.g. (01b & 10b == 00b).
&&	Logical AND operator. E.g. (01b && 10b == 01b).
^	Bitwise exclusive-OR operator; sometimes used as “raised to the power of” as well, as indicated by the context in which it is used. E.g. (01b ^ 10b == 11b). E.g. (2^2 == 4).
~	Bitwise NOT operator (also known as one’s complement). E.g. (~10b == 01b).
!	Logical NOT operator. E.g. (!10b == 00b).
==	Logical “is equal to” operator.
!=	Logical “is not equal to” operator.
<=	Less than or equal operator.
>=	Greater than or equal operator.
*	Arithmetic multiplication operator.
/	Arithmetic division operator.
<<	Shift left first operand by the number of bits specified by the 2nd operand. E.g. (01b << 01b == 10b).
>>	Shift right first operand by the number of bits specified by the 2nd operand. E.g. (10b >> 01b == 01b).
?:	Conditional operator. E.g. <i>condition ? value if true : value if false</i> . Equivalent to IF <i>condition</i> THEN <i>value if true</i> ELSE <i>value if false</i> .

Table 2: Functions

Function	Definition
ABS	ABS(integer-expression): Remove sign from signed value.
FLOOR	FLOOR(integer-expression): Rounds real number down to nearest integer.
CEIL	CEIL(real-expression): Rounds real number up to nearest integer.
MIN	MIN(integer-expression-list): Picks minimum integer or real value of comma separated list.
MAX	MAX(integer-expression-list): Picks maximum integer or real value of comma separated list.

Table 2: Functions

Function	Definition
COUNT	COUNT(integer-expression): Returns the number of binary 1's in the integer.
ROUND	ROUND(real-expression): Rounds to the nearest integer; halfway rounds away from zero.
UNIT	UNIT(fieldName UnitOfMeasure): Input operand is a register field name that defines all values with the same unit of measure. Returns the value expressed in the unit of measure for the current value of the register field. E.g. If (D18F2x210_dct[1:0]_nbp[3:0][RdPtrInit]==0101b) then (UNIT(D18F2x210_dct[1:0]_nbp[3:0][RdPtrInit] MEMCLK)==1.5).

The order in which logical operators are applied is: ~ first, & second, and | last.

For example, the equation:

$$\text{Output}[3:0] = \{A[1:0], B[3:2]\} \& C[3:0] \mid \sim D[3:0] \& E[9:6],$$

is translated as:

$$\text{Output}[3] = (A[1] \& C[3]) \mid (\sim D[3] \& E[9]);$$

$$\text{Output}[2] = (A[0] \& C[2]) \mid (\sim D[2] \& E[8]);$$

$$\text{Output}[1] = (B[3] \& C[1]) \mid (\sim D[1] \& E[7]);$$

$$\text{Output}[0] = (B[2] \& C[0]) \mid (\sim D[0] \& E[6]);$$

1.4 Definitions

Table 3: Definitions

Term	Definition
AP	Application processor. See 2.3 [Processor Initialization] .
Battery-Power	The system is running from a battery power source or otherwise undocked from a continuous power supply. Setting using this definition may be required to change during runtime.
BCS	Base configuration space. See 2.7 [Configuration Space] .
BERT	Bit error rate tester. A piece of test equipment that generates arbitrary test patterns and checks that a device under test returns them without errors.
BIST	Built-in self-test. Hardware within the processor that generates test patterns and verifies that they are stored correctly (in the case of memories) or received without error (in the case of links).
Boot VID	Boot voltage ID. This is the VDD and VDDNB voltage level that the processor requests from the external voltage regulator during the initial phase of the cold boot sequence. See 2.5.1.2 [Internal VID Registers and Encodings] .
BSC	Boot strap core. Core 0 of the BSP. Specified by MSR0000_001B[BSC].
BSP	Boot strap processor. See 2.3 [Processor Initialization] .
CAR	Use of the L2 cache as RAM during boot. See 2.3.3 [Using L2 Cache as General Storage During Boot] .
C-states	These are ACPI-defined core power states. C0 is operational. All other C-states are low-power states in which the processor is not executing code. See 2.5.3.2 [Core C-states] .
Canonical address	An address in which the state of the most-significant implemented bit is duplicated in all the remaining higher-order bits, up to bit 63.
Channel	See DRAM channel.

Table 3: Definitions

Term	Definition
Channel interleaved mode	Mode in which DRAM address space is interleaved between DRAM channels. See 2.9.6 [Memory Interleaving Modes].
CMP	Chip multi-processing. Refers to processors that include multiple cores. See 2.1 [Processor Overview].
COF	Current operating frequency of a given clock domain. See 2.5.3 [CPU Power Management].
Cold reset	PWROK is deasserted and RESET_L is asserted. See 2.3 [Processor Initialization].
Compute Unit	Two Cores that share IC, DE, FP and L2 resources. See 2.1 [Processor Overview].
Core	The instruction execution unit of the processor. See 2.1 [Processor Overview].
CPB	Core performance boost. See 2.5.3.1.1 [Application Power Management (APM)].
CpuCore-Num	Specifies the core number. See 2.4.3 [Processor Cores and Downcoring].
CPUID function X	Refers to the CPUID instruction when EAX is preloaded with X. See 3.17 [CPUID Instruction Registers].
CS	Chip select. See D18F2x[5C:40]_dct[1:0] [DRAM CS Base Address].
DCT	DRAM controller. See 2.9 [DRAM Controllers (DCTs)].
DCQ	DRAM controller queue.
DDR3	DDR3 memory technology. See 2.9 [DRAM Controllers (DCTs)].
DID	Divisor identifier. Specifies the post-PLL divisor used to reduce the COF. See 2.5.3 [CPU Power Management].
Doubleword	A 32-bit value.
Downcoring	Removal of cores. See 2.4.3 [Processor Cores and Downcoring].
DRAM channel	The part of the DRAM interface that connects to a 64-bit DIMM. For example, a processor with a 128-bit DRAM interface is said to support two DRAM channels. See 2.9 [DRAM Controllers (DCTs)].
Dual-Plane	Refers to a processor or systemboard where VDD and VDDNB are separate and may operate at independent voltage levels. Refer to 2.5.1 [Processor Power Planes And Voltage Control].
DW	Doubleword. A 32-bit value.
ECS	Extended configuration space. See 2.7 [Configuration Space].
EDS	Electrical data sheet. See 1.2 [Reference Documents].
FCH	Fusion Controller Hub. The platform device that contains the bridge to the system BIOS.
FDS	Functional data sheet; there is one FDS for each package type.
FID	Frequency identifier. Specifies the PLL frequency multiplier for a given clock domain. See 2.5.3 [CPU Power Management].
FreeR-unSample-Timer	An internal free running timer used by many power management features. The timer increments at the rate specified by D18F4x110[CSampleTimer].
GB	Gbyte or Gigabyte; 1,073,741,824 bytes.
#GP	A general-protection exception. Always produces an error code of 0. #GP(0)
GpuEnabled	GpuEnabled = (D1F0x00 != FFFF_FFFFh).
GT/s	Giga-transfers per second.

Table 3: Definitions

Term	Definition
HTC	Hardware thermal control. See 2.10.4.1 [PROCHOT_L and Hardware Thermal Control (HTC)].
HTC-active state	Hardware-controlled lower-power, lower-performance state used to reduce temperature. See 2.10.4.1 [PROCHOT_L and Hardware Thermal Control (HTC)].
IBS	Instruction based sampling. See 2.6 [Performance Monitoring].
IFCM	Isochronous flow-control mode, as defined in the link specification.
ILM	Internal loopback mode. Mode in which the link receive lanes are connected directly to the transmit lanes of the same link for testing and characterization. See D18F0x170 [Link Extended Control].
IO configuration	Access to configuration space through IO ports CF8h and CFCh. See 2.7 [Configuration Space].
IORR	IO range register. See MSRC001_00[18,16] [IO Range Base (IORR_BASE[1:0])].
IOMMU	I/O Memory Management Unit. Also known as AMD I/O Virtualization Technology. See 2.12 [IOMMU].
KB	Kbyte or Kilobyte; 1024 bytes.
L1 cache	The level 1 caches (instruction cache and the data cache) and the level 2 caches. See 2.1 [Processor Overview].
L2 cache	
Linear (virtual) address	The address generated by a core after the segment is applied.
Link	Generic term that refers to a refer to PCIe® link.
LINT	Local interrupt.
Logical address	The address generated by a core before the segment is applied.
LVT	Local vector table. A collection of APIC registers that define interrupts for local events. E.g., APIC[530:500] [Extended Interrupt [3:0] Local Vector Table].
Master abort	This is a PCI-defined term that is applied to transactions on other than PCI buses. It indicates that the transaction is terminated without affecting the intended target; reads return all 1's; writes are discarded; the master abort error code is returned in the response, if applicable; master abort error bits are set if applicable.
MB	Megabyte; 1024 KB.
MCT	Memory controller. See 2.8 [Northbridge (NB)].
MCQ	Memory controller queue. See 2.8 [Northbridge (NB)].
Micro-op	Micro-op. Instructions have variable-length encoding and many perform multiple primitive operations. The processor does not execute these complex instructions directly, but, instead, decodes them internally into simpler fixed-length instructions called macro-ops. Processor schedulers subsequently break down macro-ops into sequences of even simpler instructions called micro-ops, each of which specifies a single primitive operation. See <i>Software Optimization Guide for AMD Family 15h Processors</i> .
MEMCLK	Refers to the clock signals, M[B, A][3:0]_CLK, that are driven from the processor to DDR DIMMs.
MMIO	Memory-mapped input-output range. This is physical address space that is mapped to the IO functions such as the IO links or MMIO configuration. The IO link MMIO ranges are specified by D18F1x[1CC:180,BC:80] [MMIO Base/Limit].

Table 3: Definitions

Term	Definition
MMIO configuration	Access to configuration space through memory space. See 2.7 [Configuration Space].
MSR	Model-specific register. The core includes several MSRs for general configuration and control. See 3.18 [MSRs - MSR0000_xxxx] for the beginning of the MSR register definitions.
MTRR	Memory-type range register. The MTRRs specify the type of memory associated with various memory ranges. See MSR0000_00FE, MSR0000_020[F:0], MSR0000_02[6F:68,59:58,50], and MSR0000_02FF.
NB	Northbridge. The transaction routing block of the node. See 2.1 [Processor Overview].
NBC	Node Base Core. The lowest numbered core in the node.
NCLK	The main northbridge clock. The NCLK frequency is the NB COF.
Node ID	The identifier assigned to each node, D18F0x60[NodeId].
Node	See 2.1 [Processor Overview].
Normalized address	Addresses used by DCTs. See 2.8 [Northbridge (NB)].
OW	Octword. An 128-bit value.
ODM	On-DIMM mirroring. See D18F2x[5C:40]_dct[1:0][OnDimmMirror].
ODT	On-die termination, which is applied DRAM interface signals.
ODTS	DRAM On-die thermal sensor.
Operational frequency	The frequency at which the processor operates. See 2.5 [Power Management].
PCIe[®]	PCI Express.
PDS	Product data sheet.
Physical address	Addresses used by cores in transactions sent to the NB.
PRBS	Pseudo-random bit sequence.
Processor	See 2.1 [Processor Overview].
PSI	Power Status Indicator. See 2.5.1.3.1 [PSIx_L Bit].
P-state	Performance state. See 2.5 [Power Management].
PTE	Page table entry.
QW	Quadword. A 64-bit value.
RAS	Reliability, availability and serviceability (industry term). See 2.15 [RAS Features].
RDQ	Read data queue.
RX	Receiver.
Shutdown	A state in which the affected core waits for either INIT, RESET, or NMI. When shutdown state is entered, a shutdown special cycle is sent on the IO links.
Single-Plane	Refers to a processor or systemboard where VDD and VDDNB are tied together and operate at the same voltage level. Refer to 2.5.1 [Processor Power Planes And Voltage Control].
Slam	Refers to change the voltage to a new value in one step (as opposed to stepping). See 2.5.1.4.1 [Hardware-Initiated Voltage Transitions].

Table 3: Definitions

Term	Definition
SMAF	System management action field. This is the code passed from the SMC to the processors in STPCLK assertion messages. The action taken by the processors in response to this message is specified by D18F3x[84:80] [ACPI Power State Control].
SMBus	System management bus. Refers to the protocol on which the serial VID interface (SVI) commands and SBI are based. See 2.5.1 [Processor Power Planes And Voltage Control] and 1.2 [Reference Documents].
SMC	System management controller. This is the platform device that communicates system management state information to the processor through an IO link, typically the system IO hub.
SMI	System management interrupt. See 2.4.8.2.1 [SMM Overview].
SMM	System management mode. See 2.4.8.2 [System Management Mode (SMM)].
Speculative event	A performance monitor event counter that counts all occurrences of the event even if the event occurs during speculative code execution.
SVI2	Serial VID 2.0 interface. See 2.5.1.1 [Serial VID Interface]
SVM	Secure virtual machine. See 2.4.9 [Secure Virtual Machine Mode (SVM)].
Sync flood	The propagation of continuous sync packets to all links. This is used to quickly stop the transmission of potentially bad data when there are no other means to do so. See the link specification for additional information.
TCC	Temperature calculation circuit. See 2.10 [Thermal Functions].
Tctl	Processor temperature control value. See 2.10.4 [Temperature-Driven Logic].
TDP	Thermal design power. A power consumption parameter that is used in conjunction with thermal specifications to design appropriate cooling solutions for the processor.
Token	A scheduler entry used in various northbridge queues to track outstanding requests. See D18F3x140 [SRI to XCS Token Count] on Page 461.
TX	Transmitter.
UI	Unit interval. This is the amount of time equal to one half of a clock cycle.
UMA	Unified memory architecture. This is a type of display device that uses a frame buffer located in main memory.
UMI	Unified Media Interface. The link between the processor and the FCH.
VDD	Main power supply to the processor core logic.
VDDNB	Main power supply to the processor NB logic.
VID	Voltage level identifier. See 2.5.1 [Processor Power Planes And Voltage Control].
Virtual CAS	The clock in which CAS is asserted for the burst, N, plus the burst length (in MEMCLKs), minus 1; so the last clock of virtual CAS = $N + (BL/2) - 1$.
VRM	Voltage regulator module.
W	Word. A 16-bit value.
Warm reset	RESET_L is asserted only (while PWROK stays high). See 2.3 [Processor Initialization].
WDT	Watchdog timer. A timer that detects activity and triggers an error if a specified period of time expires without the activity. For example, see the NB watchdog timer in D18F3x40 [MCA NB Control].
WDQ	Write data queue.
XBAR	Cross bar; command packet switch. See 2.8 [Northbridge (NB)].

1.5 Changes Between Revisions and Product Variations

1.5.1 Revision Conventions

The processor revision is specified by [CPUID Fn0000_0001_EAX \[Family, Model, Stepping Identifiers\]](#) or [CPUID Fn8000_0001_EAX \[Family, Model, Stepping Identifiers\]](#). This document uses a revision letter instead of specific model numbers. The following table contains the definitions based on model and stepping used in this document. Where applicable, the processor stepping is indicated after the revision letter. All behavior marked with a revision letter apply to future revisions unless they are superseded by a change in a later revision. See the *Revision Guide for AMD Family 15h Models 10h-1Fh Processors* for additional information about revision determination.

Table 4: Processor revision conventions

Term	Definition
Revision	Revision = { CPUID Fn0000_0001_EAX[BaseModel] , CPUID Fn0000_0001_EAX[Stepping] }.
TN-A0	TN-A0 = {0, 0}.
TN-A1	TN-A1 = {0, 1}.
RL-A1	RL-A1 = {3, 1}.

1.5.2 Major Changes Relative to Family 15h Models 00h-0Fh Processors

- CPU core changes:
 - Instruction Set Architecture changes:
 - Added FMA instruction support. See [CPUID Fn0000_0001_ECX\[FMA\]](#).
 - Added F16C support. See [CPUID Fn0000_0001_ECX\[F16C\]](#).
 - Added BMI1 and TBM instruction support. See [CPUID Fn0000_0007_EBX_x0\[BMI1\]](#) and [CPUID Fn8000_0001_ECX\[TBM\]](#).
 - Added TCE support. See [CPUID Fn8000_0001_ECX\[TCE\]](#).
 - No L3 cache.
 - Other core changes:
 - Increased L1 DTLB size to 64. See [CPUID Fn8000_0005_EAX](#), [CPUID Fn8000_0005_EAX](#), and [CPUID Fn8000_0019_EAX](#).
 - Added read-only effective frequency registers. See [CPUID Fn8000_0007_EDX\[EffFreqRO\]](#).
- Memory controller (MCT) and DRAM controllers (DCTs) additions:
 - Low-voltage DDR3 support; Added 1.25V in addition to 1.35V.
 - Support for 1066 MHz (2133 MT/s) MEMCLK frequency.
 - Support for NCLK:MEMCLK frequency ratio down to 1.25:1 instead of 2:1.
- Links and IO additions:
 - PCIe[®] with generation 2 link support.
 - Per-lane power gating.
- RAS-related additions:
- General Northbridge additions:
 - NB power gating.
 - Integrated root complex.
 - Integrated graphics processor.
 - Integrated IOMMUv2.
- Power management:
 - Package C6 support.

- Bidirectional APM.
- SVI 2.0 infrastructure.
- DRAM Power Management:
 - Memory P-states.

1.5.2.1 Major Changes to Core/NB Performance Counters

- Core performance counters:
 - Added [PMCx032](#) [Misaligned Stores].
 - Added [PMCx034](#) [FP Load Buffer Stall].
 - Updated [PMCx052](#) [Ineffective Software Prefetches].
 - Updated [PMCx0D](#)[F:C].
 - Added [PMCx1C0](#) [Retired x87 Floating Point Operations].

1.5.3 Changes For Revision RL-A1

- Changes that may result in BIOS modifications.
 - Added [2.5.9.1](#) [Hybrid Boost]
 - Added [D0F0xBC_x1F428](#)[HybridBoostEn]
- Other changes
 - Added [D0F0xBC_x1F8EC](#)

2 Functional Description

2.1 Processor Overview

The *processor* is a package that contains a node consisting of (1) one to four cores (one or two compute units), (2) one PCIe® root complex with generation 2 link support, (3) two 64-bit DDR3 interfaces for communication to system memory, and (4) one communication packet routing block referred to as the *northbridge* (NB).

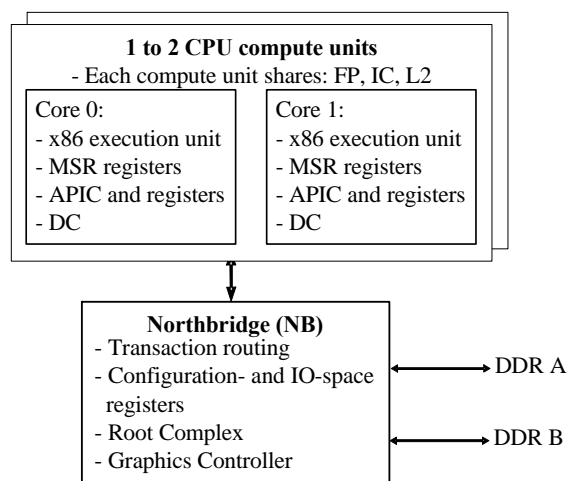


Figure 1: A processor

Each *compute unit* includes 2 cores each having an x86 instruction execution logic and first-level (L1) data cache. The FP unit, second level (L2) general-purpose cache, and first-level instruction cache are shared between the cores of a compute unit. There is a set of MSRs and APIC registers associated with each core. Processors that include multiple cores are said to incorporate *chip multi-processing* or CMP.

The links are input-output links, as defined by the PCI Express Base Specification.

Each DRAM interface supports a 64-bit DDR3 memory channel.

The NB routes transactions between the cores, the link, and the DRAM interfaces. It includes the configuration register space for the device.

2.2 System Overview

The following diagram illustrates the expected system architecture.

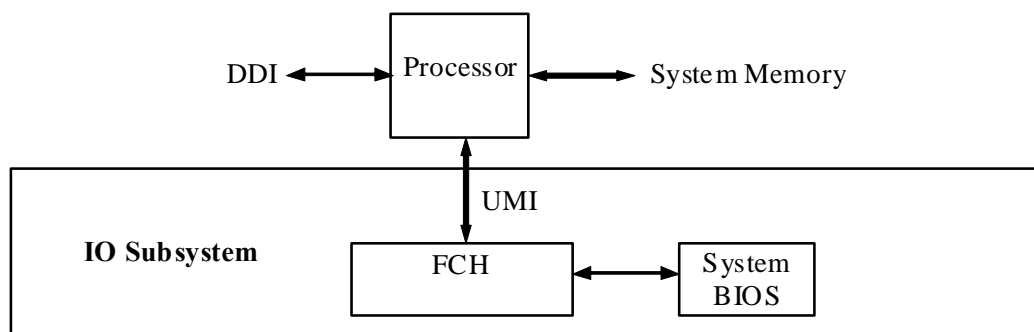


Figure 2: System Diagram

2.3 Processor Initialization

This section describes the initialization sequence after a cold reset.

Core 0 of the processor, the bootstrap core (BSC), begins executing code from the reset vector. The remaining cores do not fetch code until their enable bits are set ([D18F0x1DC\[CpuEn\]](#)).

2.3.1 BSC initialization

The BSC must perform the following tasks as part of boot.

- Store BIST information from the EAX register into an unused processor register.
- [D18F0x6C\[InitDet\]](#) may be used by BIOS to differentiate between INIT and cold/warm reset.
- Determine type of startup using [D18F0x6C\[ColdRstDet\]](#).
 - If this is a warm reset then BIOS may check for valid MCA errors and if present save the status for later use. See [2.15.1.6 \[Handling Machine Check Exceptions\]](#).
- Enable the cache, program the MTRRs for CAR and initialize CAR. See [2.3.3 \[Using L2 Cache as General Storage During Boot\]](#).
- Setup the SMU.
- Setup of APIC ([2.4.8.1.3 \[ApicId Enumeration Requirements\]](#)).
- Setup the link configuration [2.11.3.2 \[Link Configurations\]](#).
- Setup the root complex and initialize the I/O links [2.11.4.2 \[Link Configuration and Initialization\]](#).
- If required, reallocate data and flow control buffers of the links (see [D18F0x90 \[Upstream Base Channel Buffer Count\]](#) and [D18F0x94 \[Link Isochronous Channel Buffer Count\]](#)) and issue system warm reset.
- Configure the DRAM controllers.
- Configure processor power management. See [2.5 \[Power Management\]](#).
- If supported, allow other cores to begin fetching instructions by setting [D18F0x1DC\[CpuEn\]](#) in the PCI configuration space of the nodes. See [2.4.3 \[Processor Cores and Downcoring\]](#).

2.3.2 AP initialization

All other processor cores other than core 0 begin executing code from the reset vector. They must perform the following tasks as part of boot.

- Store BIST information from the eax register into an unused processor register.

- **D18F0x6C**[InitDet] may be used by BIOS to differentiate between INIT and cold/warm reset.
- Determine the history of this reset using the **D18F0x6C** [Link Initialization Control] [ColdRstDet] bit:
 - If this is a warm reset then BIOS may check for valid MCA errors and if present save the status for use later. See 2.15.1.6 [Handling Machine Check Exceptions].
- Set up the local APIC. See 2.4.8.1.3 [ApicId Enumeration Requirements].
- Configure processor power management. See 2.5 [Power Management].

2.3.3 Using L2 Cache as General Storage During Boot

Prior to initializing the DRAM controller for system memory, BIOS may use the L2 cache of each core as general storage.

The L2 cache as general storage is described as follows:

- Each compute unit has its own L2 cache.
- BIOS manages the mapping of the L2 storage such that cacheable accesses do not cause L2 victims.
- The L2 size, L2 associativity, and L2 line size is determined by reading **CPUID Fn8000_0006_ECX**[L2Size, L2Assoc, L2LineSize]. L2WayNum is defined to be the number of ways indicated by the L2Assoc code.
 - The L2 cache is viewed as (L2Size/L2LineSize) cache lines of storage, organized as L2WayNum ways, each way being (L2Size/L2WayNum) in size.
 - E.g. L2Assoc=8 so L2WayNum=16 (there are 16 ways). If L2Size=512KB then there are 16 blocks of cache, each 512KB/16 in size, or 32KB each.
 - For each of the following values of L2Size, the following values are defined:
 - L2Size=1 MB: L2Tag=PhysAddr[39:16], L2WayIndex=PhysAddr[15:6].
 - L2Size=2 MB: L2Tag=PhysAddr[39:17], L2WayIndex=PhysAddr[16:6].
 - PhysAddr[5:0] addresses the L2LineSize number of bytes of storage associated with the cache line.
 - The L2 cache, when allocating a line at L2WayIndex:
 - Picks an invalid way before picking a valid way.
 - Prioritizes the picking of invalid ways such that way L2WayNum-1 is the highest priority and 0 is the lowest priority.
 - It is recommended that BIOS assume a simpler allocation of L2 cache memory, being L2WayNum size-aligned blocks of memory, each being L2Size/L2WayNum bytes.
 - BIOS can rely on a minimum L2Size of 512 KB and can rely on being able to use a minimum of 14 ways for general storage. See **CPUID Fn8000_0006_ECX**[L2Size]. See initialization requirements below for **MSRC001_1023**[L2WayLock, L2FirstLockedWay].

The following memory types are supported:

- **WP-IO**: BIOS ROM may be assigned the write-protect IO memory type and may be accessed read-only as data and fetched as instructions.
 - WP-IO accesses, both read and write, do not get evicted to the L2 and therefore do not need to be considered for allocation into the L2.
- **WB-DRAM**: General storage may be assigned the write-back DRAM memory type and may be accessed as read-write data, but not accessed by instruction fetch.
 - BIOS initializes an L2LineSize sized and aligned location in the L2 cache, mapped as write-back DRAM, with 1 read to at least 1 byte of the L2LineSize sized and aligned WB-DRAM address. BIOS may store to a line only after it has been allocated by a load.
 - Fills, sent to the disabled memory controller, return undefined data.
- All of memory space that is not accessed as WP-IO or WB-DRAM space must be marked as UC memory type.
- In order to prevent victimizing L2 data, no more than L2WayNum cache lines accessed as WB-DRAM may have the same L2WayIndex.

- Software does not need to know which ways the L2WayNum lines are allocated to for any given value of L2WayIndex, only that invalid ways will be selected for allocation before valid ways will be selected for allocation.
- Software is not allowed to deallocate a line in the L2 by using CLFLUSH. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Performance monitor event `PMCx07F[1]`, titled “L2 Writebacks to system“, can be used to indicate whether L2 dirty data was lost by being victimized and sent to the disabled memory controller.

The following requirements must be satisfied prior to using the cache as general storage:

- Paging must be disabled.
- `MSRC001_0015[INVDWBINVD]=0`.
- `MSRC001_1020[DisSS]=1`.
- `MSRC001_1021[DisSpecTlbRld]=1`. Disable speculative ITLB reloads.
- `MSRC001_1022[DisSpecTlbRld]=1`. Disable speculative DTLB reloads.
- `MSRC001_1022[DisHwPf]=1`.
- `MSRC001_102B[CombineCr0Cd]=0`.
- CLFLUSH, INVD, and WBINVD must not be used during CAR but may be used when tearing down CAR for all compute units on a node.
- The BIOS must not use SSE, or MMX™ instructions, with the exception of the following list: MOVD, MOVQ, MOVDQA, MOVQ2DQ, MOVDQ2Q.
- The BIOS must not enable exceptions, page-faults, and other interrupts.
- BIOS must not use software prefetches.
- UC-DRAM: All of DRAM that is not accessed as WB-DRAM space must be marked as UC memory type.
- If (`MSRC001_1023[L2WayLock]==1`) then:
 - Only the ways 0 through (`MSRC001_1023[L2FirstLockedWay]-1`) may be used for general storage.
 - BIOS can rely on `MSRC001_1023[L2FirstLockedWay]` to have a minimum value of Eh.
- If (`MSRC001_1023[L2WayLock]==0`) then:
 - Set `MSRC001_1023[L2WayLock]=1`.
 - Set `MSRC001_1023[L2FirstLockedWay]=Fh`.

When BIOS has completed using the cache for general storage the following steps are followed:

1. An INVD instruction should be executed on each core that used cache as general storage; an INVD should only be issued when all cores on all nodes have completed using the cache for general storage.
2. If DRAM is initialized and there is data in the cache that needs to get moved to main memory, CLFLUSH or WBINVD may be used instead of INVD, but software must ensure that needed data in main memory is not overwritten.
3. Program the following configuration state:
 - `MSRC001_0015[INVDWBINVD]=1`.
 - `MSRC001_1020[DisSS]=0`.
 - `MSRC001_1021[DisSpecTlbRld]=0`.
 - `MSRC001_1022[DisSpecTlbRld]=0`.
 - `MSRC001_1022[DisHwPf]=0`.
 - If (`(MSRC001_1023[L2WayLock]==1) & (MSRC001_1023[L2FirstLockedWay]==Fh)`), program `MSRC001_1023[L2WayLock]=0`.

2.4 Core

The majority of the behavioral definition of the core is specified in the AMD64 Architecture Programmer's Manual. See 1.2 [Reference Documents].

2.4.1 Compute Unit

Unless otherwise specified the processor configuration interface hides the compute unit implementation and presents software with homogenous cores, each independent of the other.

Software may use [D18F5x80](#)[Enabled, DualCore] in order to associate a core with a compute unit. This information can be useful because some configuration settings are determined based on active compute units and core performance may vary based on resource sharing within a compute unit.

2.4.1.1 Registers Shared by Cores in a Compute Unit

Some MSR's are implemented one copy per compute unit instead of per core; these MSR's are designated as [SharedC](#) (shared coherent) or [SharedNC](#) (shared non-coherent). The absence of [SharedC/SharedNC](#) implies not-shared, which is the normal per-core instance programming model for RDMSR/WRMSR.

Programming rules for [SharedC](#) and [SharedNC](#) registers:

- Software must ensure that a shared MSR written by one core on a compute unit will not cause a problem for software that is running on the other core of the compute unit.
- [SharedC](#): A write to a [SharedC](#) MSR does not have to be written to the other core of the compute unit in order for the other core to see the updated value.
- [SharedNC](#): A write to a [SharedNC](#) MSR has to be written to both cores of the compute unit in order for both cores to see the updated value.
 - If software can know that the other core has not read the [SharedNC](#) MSR since the last warm reset, then a write is not needed to the [SharedNC](#) MSR on the other core.
 - Software may not rely on the other core maintaining the previous value of the [SharedNC](#) MSR.
 - The [SharedNC](#) MSR's are: [MSRC001_00](#)[35:30], [MSRC001_00](#)[53:50], [MSRC001_0054](#), [MSRC001_0055](#).
- A read-modify-write of a shared MSR register is not atomic. Software must ensure atomicity between the cores that could simultaneously read-modify-write the shared register.

2.4.2 Virtual Address Space

The processor supports 48 address bits of virtual memory space (256 TB) as indicated by [CPUID Fn8000_0008_EAX](#).

2.4.3 Processor Cores and Downcoring

The processor supports downcoring as follows:

- The cores of a compute unit may be software downcored by [D18F3x190](#)[DisCore]. See 2.4.3.1 [Software Downcoring using [D18F3x190](#)[DisCore]].
 - Both cores of a compute unit must be downcored if either core needs to be downcored.
 - The cores of a compute unit are even core x and odd core x+1..
 - Clocks are turned off and power is gated to downcored compute units. The power savings is the same as CC6.
 - There must be at least 1 compute unit enabled.
 - Downcoring affects: [D18F5x84](#)[CmpCap], [CPUID Fn8000_0008_ECX](#)[NC].
- An implemented (physical) core that is downcored is not visible to software. Cores that are not downcored

are numbered logically in a contiguous manner.

- [D18F5x80 \[Compute Unit Status\]](#) reports core topology information to software.
- The number of cores specified in [CPUID Fn8000_0008_ECX\[NC\]](#) must be the same as the number of cores enabled in [D18F0x1DC\[CpuEn\]](#).
- The core number, *CpuCoreNum*, is provided to SW running on each core through [CPUID Fn0000_0001_EBX\[LocalApicId\]](#) and [APIC20\[ApicId\]](#); *CpuCoreNum* also affects [D18F0x1DC\[CpuEn\]](#). *CpuCoreNum*, varies as the lowest integers from 0 to [D18F5x84\[CmpCap\]](#), based on the number of enabled cores; e.g., a 4-core node with 1 core disabled results in cores reporting *CpuCoreNum* values of 0, 1, and 2 regardless of which core is disabled. The boot core is always the core reporting *CpuCoreNum*=0.

2.4.3.1 Software Downcoring using [D18F3x190\[DisCore\]](#)

Cores may be downcored by [D18F3x190\[DisCore\]](#).

Software is required to use [D18F3x190\[DisCore\]](#) as follows:

- Once a core has been removed by [D18F3x190\[DisCore\]](#)=1, it cannot be added back without a cold reset. E.g. Software may only set *DisCore* bits, never clear them.
- If the number of cores in the system is changed, then [D18F0x60\[CpuCnt\]](#) must be updated to reflect the new value after the warm reset.
- BIOS should configure [MSRC001_102A\[ThrottleNbInterface\]](#) to reflect the number of enabled compute units.

2.4.4 Physical Address Space

The core supports 48 address bits of coherent memory space (256 terabytes) as indicated by [CPUID Fn8000_0008_EAX \[Long Mode Address Size Identifiers\]](#). However the NB only supports 40 address bits. The processor master aborts the following upper-address transactions (to address *PhysAddr*):

- Link or core requests with non-zero *PhysAddr*[63:40].

2.4.5 System Address Map

The processor defines a reserved memory address region starting at 0000_00FD_0000_0000h and extending up to 0000_0100_0000_0000h. System software must not map memory into this region. Downstream host accesses to the reserved address region results in a page fault. Upstream system device accesses to the reserved address region results in an undefined operation.

2.4.5.1 Memory Access to the Physical Address Space

All memory accesses to the physical address space from a core are sent to its associated northbridge (NB). All memory accesses from a link are routed through the NB.

A core access to physical address space has two important attributes that must be determined before issuing the access to the NB: the memory type (e.g., WB, WC, UC; as described in the MTRRs) and the access destination (DRAM or MMIO).

This mechanism is managed by the BIOS and does not require any setup or changes by system software.

2.4.5.1.1 Determining Memory Type

The memory type for a core access is determined by the highest priority of the following ranges that the access falls in: 1==Lowest priority.

1. The memory type as determined by architectural mechanisms.

- See the APM2 chapter titled “Memory System”, sections “Memory-Type Range Registers” and “Page-Attribute Table Mechanism”.
 - See the APM2 chapter titled “Nested Paging”, section “Combining Memory Types, MTRRs”.
 - See [MSR0000_02FF \[MTRR Default Memory Type \(MTRRdefType\)\]](#), [MSR0000_020\[F:0\] \[Variable-Size MTRRs Base/Mask\]](#), [MSR0000_02\[6F:68,59:58,50\] \[Fixed-Size MTRRs\]](#).
2. TSeg & ASeg SMM mechanism. (see [MSRC001_0112](#) and [MSRC001_0113](#))
 3. CR0[CD]: If (CR0[CD]==1) then MemType=CD.
 4. MMIO config space, APIC space.
 - MMIO APIC space and MMIO config space must not overlap.
 - MemType=UC.
 - See [2.4.8.1.2 \[APIC Register Space\]](#) and [2.7 \[Configuration Space\]](#).
 5. If (“In SMM Mode”) && (([MSRC001_0113\[AValid\]](#) && “The address does not fall within the ASeg region”) || ([MSRC001_0113\[TValid\]](#) && “The address does not fall within the TSeg region”))) then MemType=CD.

2.4.5.1.2 Determining The Access Destination for Core Accesses

The access destination, DRAM or MMIO, is based on the highest priority of the following ranges that the access falls in: 1==Lowest priority.

1. RdDram/WrDram as determined by [MSRC001_001A \[Top Of Memory \(TOP_MEM\)\]](#) and [MSRC001_001D \[Top Of Memory 2 \(TOM2\)\]](#).
2. The IORRs. (see [MSRC001_00\[18,16\]](#) and [MSRC001_00\[19,17\]](#)).
3. The fixed MTRR’s. (see [MSR0000_02\[6F:68,59:58,50\] \[Fixed-Size MTRRs\]](#))
4. TSeg & ASeg SMM mechanism. (see [MSRC001_0112](#) and [MSRC001_0113](#))
5. MMIO config space, APIC space.
 - MMIO APIC space and MMIO config space must not overlap.
 - RdDram=IO, WrDram=IO.
 - See [2.4.8.1.2 \[APIC Register Space\]](#) and [2.7 \[Configuration Space\]](#).
6. NB address space routing. See [2.8.2.1.1 \[DRAM and MMIO Memory Space\]](#).

2.4.6 Timers

Each core includes the following timers. These timers do not vary in frequency regardless of the current P-state or C-state.

- [MSR0000_0010 \[Time Stamp Counter \(TSC\)\]](#); the TSC increments at the rate specified by the P0 P-state; see [2.5.3.1.2.1 \[Software P-state Numbering\]](#) and [MSRC001_00\[6B:64\] \[P-state \[7:0\]\]](#).
- The APIC timer ([APIC380](#) and [APIC390](#)), which increments at the rate of 2xCLKIN; the APIC timer may increment in units of between 1 and 8.

2.4.7 Implicit Conditions for TLB Invalidation

The following family specific conditions will cause all TLB’s for both cores of the compute unit to be invalidated; except [MSR0000_0277](#) which will only clear the TLB’s for the core that did the MSR write. The architectural conditions that cause TLB invalidation are documented by the APM2 section titled “Translation-Lookaside Buffer (TLB)”; see “Implicit Invalidations”.

- [MSR0000_020\[F:0\] \[Variable-Size MTRRs Base/Mask\]](#)
- [MSR0000_02\[6F:68,59:58,50\] \[Fixed-Size MTRRs\]](#)
- [MSR0000_0277 \[Page Attribute Table \(PAT\)\]](#) (TLB’s not cleared for the other core)
- [MSR0000_02FF \[MTRR Default Memory Type \(MTRRdefType\)\]](#)
- [MSRC001_0010 \[System Configuration \(SYS_CFG\)\]](#) write.

- [MSRC001_00\[18,16\]](#) [IO Range Base (IORR_BASE[1:0])] write.
- [MSRC001_00\[19,17\]](#) [IO Range Mask (IORR_MASK[1:0])] write.
- [MSRC001_001A](#) [Top Of Memory (TOP_MEM)] write.
- [MSRC001_001D](#) [Top Of Memory 2 (TOM2)] write.
- [MSRC001_1023](#) [Combined Unit Configuration (CU_CFG)] write.
- [MSRC001_102A](#) [Combined Unit Configuration 2 (CU_CFG2)] write.

2.4.8 Interrupts

2.4.8.1 Local APIC

The local APIC contains logic to receive interrupts from a variety of sources and to send interrupts to other local APICs, as well as registers to control its behavior and report status. Interrupts can be received from:

- IO devices including the IO hub (IO APICs)
- Other local APICs (inter-processor interrupts)
- APIC timer
- Thermal events
- Performance counters
- Legacy local interrupts from the IO hub (INTR and NMI)
- APIC internal errors

The APIC timer, thermal events, performance counters, local interrupts, and internal errors are all considered local interrupt sources, and their routing is controlled by local vector table entries. These entries assign a message type and vector to each interrupt, allow them to be masked, and track the status of the interrupt.

IO and inter-processor interrupts have their message type and vector assigned at the source and are unaltered by the local APIC. They carry a destination field and a mode bit that together determine which local APIC(s) accepts them. The destination mode (DM) bit specifies if the interrupt request packet should be handled in physical or logical destination mode. If the destination field is FFh, the interrupt is broadcast and is accepted by all local APICs regardless of destination mode. If the destination field matches the broadcast value of FFh, then the interrupt is a broadcast interrupt and is accepted by all local APICs regardless of destination mode.

2.4.8.1.1 Detecting and Enabling

APIC is detected and enabled via [CPUID Fn0000_0001_EDX](#)[APIC].

The local APIC is enabled via [MSR0000_001B](#)[ApicEn]. Reset forces APIC disabled.

2.4.8.1.2 APIC Register Space

MMIO APIC space:

- Memory mapped to a 4 KB range. The memory type of this space is forced to the UC memory type. The base address of this range is specified by {[MSR0000_001B](#)[ApicBar[47:12]], 000h}.
- The mnemonic is defined to be APICXX; XX is the byte address offset from the base address.
- MMIO APIC registers in xAPIC mode is defined by the register from [APIC20](#) to [APIC\[530:500\]](#).
- Treated as normal memory space when APIC is disabled, as specified by [MSR0000_001B](#)[ApicEn].

2.4.8.1.3 ApicId Enumeration Requirements

System hardware and BIOS must ensure that the number of cores per processor (NC) exposed to the operating system by all tables, registers, and instructions across all cores in the processor is identical. See [2.4.10.1 \[Multi-](#)

[Core Support](#)] to derive NC.

Operating systems are expected to use [CPUID Fn8000_0008_ECX\[ApicIdCoreIdSize\]](#), the number of least significant bits in the Initial APIC ID that indicate core ID within a processor, in constructing per-core CPUID masks. ([ApicIdCoreIdSize\[3:0\]](#) determines the maximum number of cores (MNC) that the processor could theoretically support, not the actual number of cores that are actually implemented or enabled on the processor, as indicated by [CPUID Fn8000_0008_ECX\[NC\]](#).) BIOS must use the [ApicId MNC](#) rule when assigning [APIC20 \[APIC ID\]\[ApicId\]](#) values as described below.

[ApicId MNC](#) rule: The [ApicId](#) of core *j* must be enumerated/assigned as:

$$\text{ApicId}[\text{core}=j] = (\text{OFFSET_IDX}) * \text{MNC} + j$$

Where [OFFSET_IDX](#) is an integer offset (0 to N) used to shift up the core [ApicId](#) values to allow room for IOAPIC devices.

It is recommended that BIOS use the following APIC ID assignments for the broadest operating system support. Given $N = \text{MNC}$ and $M = \text{Number_Of_IOAPICs}$:

- Assign the core [ApicId](#)'s first from 0 to N-1, and the IOAPIC IDs from N to N+(M-1).

2.4.8.1.4 Physical Destination Mode

The interrupt is only accepted by the local APIC whose [APIC20\[ApicId\]](#) matches the destination field of the interrupt. Physical mode allows up to 255 APICs to be addressed individually.

2.4.8.1.5 Logical Destination Mode

A local APIC accepts interrupts selected by [APICD0 \[Logical Destination \(LDR\)\]](#) and the destination field of the interrupt using either cluster or flat format as configured by [APICE0\[Format\]](#).

If flat destinations are in use, bits 7-0 of [APICD0\[Destination\]](#) are checked against bits 7-0 of the arriving interrupt's destination field. If any bit position is set in both fields, the local APIC is a valid destination. Flat format allows up to 8 APICs to be addressed individually.

If cluster destinations are in use, bits 7-4 of [APICD0\[Destination\]](#) are checked against bits 7-4 of the arriving interrupt's destination field to identify the cluster. If all of bits 7-4 match, then bits 3-0 of [APICD0\[Destination\]](#) and the interrupt destination are checked for any bit positions that are set in both fields to identify processors within the cluster. If both conditions are met, the local APIC is a valid destination. Cluster format allows 15 clusters of 4 APICs each to be addressed.

2.4.8.1.6 Interrupt Delivery

SMI, NMI, INIT, Startup, and External interrupts are classified as non-vectorized interrupts.

When an APIC accepts a non-vectorized interrupt, it is handled directly by the processor instead of being queued in the APIC. When an APIC accepts a fixed or lowest-priority interrupt, it sets the bit in [APIC\[270:200\] \[Interrupt Request \(IRR\)\]](#) corresponding to the vector in the interrupt. For local interrupt sources, this comes from the vector field in that interrupt's local vector table entry. The corresponding bit in [APIC\[1F0:180\] \[Trigger Mode \(TMR\)\]](#) is set if the interrupt is level-triggered and cleared if edge-triggered. If a subsequent interrupt with the same vector arrives when the corresponding bit in [APIC\[270:200\]\[RequestBits\]](#) is already set, the two

interrupts are collapsed into one. Vectors 15-0 are reserved.

2.4.8.1.7 Vectored Interrupt Handling

APIC80 [Task Priority (TPR)] and **APICA0 [Processor Priority (PPR)]** each contain an 8-bit priority divided into a main priority (bits 7-4) and a priority sub-class (bits 3-0). The task priority is assigned by software to set a threshold priority at which the processor is interrupted.

The processor priority is calculated by comparing the main priority (bits 7-4) of **APIC80[Priority]** to bits 7-4 of the 8-bit encoded value of the highest bit set in **APIC[170:100] [In-Service (ISR)]**. The processor priority is the higher of the two main priorities.

The processor priority is used to determine if any accepted interrupts (indicated by **APIC[270:200][Request-Bits]**) are high enough priority to be serviced by the processor. When the processor is ready to service an interrupt, the highest bit in **APIC[270:200][RequestBits]** is cleared, and the corresponding bit is set in **APIC[170:100][InServiceBits]**.

When the processor has completed service for an interrupt, it performs a write to **APICB0 [End of Interrupt]**, clearing the highest bit in **APIC[170:100][InServiceBits]** and causing the next-highest interrupt to be serviced. If the corresponding bit in **APIC[1F0:180][TriggerModeBits]** is set, a write to **APICB0** is performed on all APICs to complete service of the interrupt at the source.

2.4.8.1.8 Interrupt Masking

Interrupt masking is controlled by the **APIC410 [Extended APIC Control]**. If **APIC410[IerCap]** is set, **APIC[4F0:480] [Interrupt Enable]** are used to mask interrupts. Any bit in **APIC[4F0:480][InterruptEnableBits]** that is clear indicates the corresponding interrupt is masked. A masked interrupt is not serviced and the corresponding bit in **APIC[270:200][RequestBits]** remains set.

2.4.8.1.9 Spurious Interrupts

In the event that the task priority is set to or above the level of the interrupt to be serviced, the local APIC delivers a spurious interrupt vector to the processor, as specified by **APICF0 [Spurious-Interrupt Vector (SVR)]**. **APIC[170:100]** is not changed and no write to **APICB0** occurs.

2.4.8.1.10 Spurious Interrupts Caused by Timer Tick Interrupt

A typical interrupt is asserted until it is serviced. An interrupt is deasserted when software clears the interrupt status bit within the interrupt service routine. Timer tick interrupt is an exception, since it is deasserted regardless of whether it is serviced or not.

The processor is not always able to service interrupts immediately (i.e. when interrupts are masked by clearing **EFLAGS.IM**).

If the processor is not able to service the timer tick interrupt for an extended period of time, the **INTR** caused by the first timer tick interrupt asserted during that time is delivered to the local APIC in **ExtInt** mode and latched, and the subsequent timer tick interrupts are lost. The following cases are possible when the processor is ready to service interrupts:

- An **ExtInt** interrupt is pending, and **INTR** is asserted. This results in timer tick interrupt servicing. This occurs 50 percent of the time.
- An **ExtInt** interrupt is pending, and **INTR** is deasserted. The processor sends the interrupt acknowledge

cycle, but when the PIC receives it, INTR is deasserted, and the PIC sends a spurious interrupt vector. This occurs 50 percent of the time.

There is a 50 percent probability of spurious interrupts to the processor.

2.4.8.1.11 Lowest-Priority Interrupt Arbitration

Fixed, remote read, and non-vectored interrupts are accepted by their destination APICs without arbitration.

Delivery of lowest-priority interrupts requires all APICs to arbitrate to determine which one accepts the interrupt. If `APICF0[FocusDisable]` is clear, then the focus processor for an interrupt always accepts the interrupt. A processor is the focus of an interrupt if it is already servicing that interrupt (corresponding bit in `APIC[170:100][InServiceBits]` is set) or if it already has a pending request for that interrupt (corresponding bit in `APIC[270:200][RequestBits]` is set). If `APIC410[IerCap]` is set the interrupt must also be enabled in `APIC[4F0:480][InterruptEnableBits]` for a processor to be the focus processor. If there is no focus processor for an interrupt, or focus processor checking is disabled, then each APIC calculates an arbitration priority value, stored in `APIC90 [Arbitration Priority (APR)]`, and the one with the lowest result accepts the interrupt.

The arbitration priority value is calculated by comparing `APIC80[Priority]` with the 8-bit encoded value of the highest bit set in `APIC[270:200][RequestBits]` (IRRVec) and the 8-bit encoded value of the highest bit set `APIC[170:100][InServiceBits]` (ISRVec). If `APIC410[IerCap]` is set the IRRVec and ISRVec are based off the highest enabled interrupt. The main priority bits 7-4 are compared as follows:

```
If (APIC80[Priority[7:4]] >= IRRVec[7:4] and APIC80[Priority[7:4]] > ISRVec[7:4])
Then APIC90[Priority] = APIC80[Priority]
Else if (IRRVec[7:4] > ISRVec[7:4]) APIC90[Priority] = {IRRVec[7:4],0h}
Else APIC90[Priority] = {ISRVec[7:4],0h}
```

2.4.8.1.12 Inter-Processor Interrupts

`APIC300 [Interrupt Command Low (ICR Low)]` and `APIC310 [Interrupt Command High (ICR High)]` provide a mechanism for generating interrupts in order to redirect an interrupt to another processor, originate an interrupt to another processor, or allow a processor to interrupt itself. A write to register `APIC300` causes an interrupt to be generated with the properties specified by the `APIC300` and `APIC310` fields.

2.4.8.1.13 APIC Timer Operation

The local APIC contains a 32-bit timer, controlled by `APIC320 [LVT Timer]`, `APIC380 [Timer Initial Count]`, and `APIC3E0 [Timer Divide Configuration]`. The processor bus clock is divided by the value in `APIC3E0[Div]` to obtain a time base for the timer. When `APIC380[Count]` is written, the value is copied into `APIC390 [Timer Current Count]`. `APIC390[Count]` is decremented at the rate of the divided clock. When the count reaches 0, a timer interrupt is generated with the vector specified in `APIC320[Vector]`. If `APIC320[Mode]` specifies periodic operation, `APIC390[Count]` is reloaded with the `APIC380[Count]` value, and it continues to decrement at the rate of the divided clock. If `APIC320[Mask]` is set, timer interrupts are not generated.

2.4.8.1.14 Generalized Local Vector Table

All LVTs (`APIC330` to `APIC3[60:50]`, and `APIC[530:500]`) support a generalized message type as follows:

- 000b=Fixed
- 010b=SMI
- 100b=NMI
- 111b=ExtINT

- All other messages types are reserved.

2.4.8.1.15 State at Reset

At power-up or reset, the APIC is hardware disabled (`MSR0000_001B[ApicEn]=0`) so only SMI, NMI, INIT, and ExtInt interrupts may be accepted.

The APIC can be software disabled through `APICF0[APICSWEn]`. The software disable has no effect when the APIC is hardware disabled.

When a processor accepts an INIT interrupt, the APIC is reset as at power-up, with the exception that:

- `APIC20[ApicId]` is unaffected.
- Pending APIC register writes complete.

2.4.8.2 System Management Mode (SMM)

System management mode (SMM) is typically used for system control activities such as power management. These activities are typically transparent to the operating system.

2.4.8.2.1 SMM Overview

SMM is entered by a core on the next instruction boundary after a system management interrupt (SMI) is received and recognized. A core may be programmed to broadcast a special cycle to the system, indicating that it is entering SMM mode. The core then saves its state into the SMM memory state save area and jumps to the SMI service routine (or SMI handler). The pointer to the SMI handler is specified by MSRs. The code and data for the SMI handler are stored in the SMM memory area, which may be isolated from the main memory accesses.

The core returns from SMM by executing the RSM instruction from the SMI handler. The core restores its state from the SMM state save area and resumes execution of the instruction following the point where it entered SMM. The core may be programmed to broadcast a special bus cycle to the system, indicating that it is exiting SMM mode.

2.4.8.2.2 Operating Mode and Default Register Values

The software environment after entering SMM has the following characteristics:

- Addressing and operation is in Real mode.
 - A far jump, call or return in the SMI handler can only address the lower 1M of memory, unless the SMI handler first switches to protected mode.
 - If (`MSRC001_0111[SmmBase]>=0010_0000h`) then:
 - The value of the CS selector is undefined upon SMM entry.
 - The undefined CS selector value should not be used as the target of a far jump, call, or return.
- 4-Gbyte segment limits.
- Default 16-bit operand, address, and stack sizes (instruction prefixes can override these defaults).
- Control transfers that do not override the default operand size truncate the EIP to 16 bits.
- Far jumps or calls cannot transfer control to a segment with a base address requiring more than 20 bits, as in Real mode segment-base addressing, unless a change is made into protected mode.
- A20M# is disabled. A20M# assertion or deassertion have no affect during SMM.
- Interrupt vectors use the Real mode interrupt vector table.
- The IF flag in EFLAGS is cleared (INTR is not recognized).
- The TF flag in EFLAGS is cleared.

- The NMI and INIT interrupts are masked.
- Debug register DR7 is cleared (debug traps are disabled).

The SMM base address is specified by [MSRC001_0111 \[SMM Base Address \(SMM_BASE\)\]](#)[SmmBase]. Important offsets to the base address pointer are:

- [MSRC001_0111](#)[SmmBase] + 8000h: SMI handler entry point.
- [MSRC001_0111](#)[SmmBase] + FE00h - FFFFh: SMM state save area.

2.4.8.2.3 SMI Sources And Delivery

The processor accepts SMIs as link-defined interrupt messages only. The core/node destination of these SMIs is a function of the destination field of these messages. However, the expectation is that all such SMI messages are specified to be delivered globally (to all cores of all nodes).

There are also several local events that can trigger SMIs. However, these local events do not generate SMIs directly. Each of them triggers a programmable IO cycle that is expected to target the SMI command port in the IO hub and trigger a global SMI interrupt message back to the coherent fabric.

Local sources of SMI events that generate the IO cycle specified in [MSRC001_0056 \[SMI Trigger IO Cycle\]](#) are:

- In the core, as specified by:
 - [MSRC001_0022 \[Machine Check Exception Redirection\]](#).
 - [MSRC001_00\[53:50\] \[IO Trap \(SMI_ON_IO_TRAP_\[3:0\]\)\]](#).
- All local APIC LVT registers programmed to generate SMIs.

The status for these is stored in [SMMFEC4](#).

2.4.8.2.4 SMM Initial State

After storing the save state, execution starts at [MSRC001_0111](#)[SmmBase] + 08000h. The SMM initial state is specified in the following table.

Table 5: SMM Initial State

Register	SMM Initial State
CS	SmmBase[19:4]
DS	0000h
ES	0000h
FS	0000h
GS	0000h
SS	0000h
General-Purpose Registers	Unmodified
EFLAGS	0000_0002h
RIP	0000_0000_0000_8000h
CR0	Bits 0, 2, 3, and 31 cleared (PE, EM, TS, and PG); remainder is unmodified
CR4	0000_0000_0000_0000h
GDTR	Unmodified
LDTR	Unmodified

Table 5: SMM Initial State

Register	SMM Initial State
IDTR	Unmodified
TR	Unmodified
DR6	Unmodified
DR7	0000_0000_0000_0400h
EFER	All bits are cleared except bit 12 (SVME) which is unmodified.

2.4.8.2.5 SMM Save State

In the following table, the offset field provides the offset from the SMM base address specified by [MSRC001_0111 \[SMM Base Address \(SMM_BASE\)\]](#).

Table 6: SMM Save State

Offset	Size	Contents		Access
FE00h	Word	ES	Selector	Read-only
FE02h	6 Bytes		Reserved	
FE08h	Quadword		Descriptor in memory format	
FE10h	Word	CS	Selector	Read-only
FE12h	6 Bytes		Reserved	
FE18h	Quadword		Descriptor in memory format	
FE20h	Word	SS	Selector	Read-only
FE22h	6 Bytes		Reserved	
FE28h	Quadword		Descriptor in memory format	
FE30h	Word	DS	Selector	Read-only
FE32h	6 Bytes		Reserved	
FE38h	Quadword		Descriptor in memory format	
FE40h	Word	FS	Selector	Read-only
FE42h	2 Bytes		Reserved	
FE44h	Doubleword		FS Base {16'b[47], 47:32} ¹	
FE48h	Quadword		Descriptor in memory format	
FE50h	Word	GS	Selector	Read-only
FE52h	2 Bytes		Reserved	
FE54h	Doubleword		GS Base {16'b[47], 47:32} ¹	
FE58h	Quadword		Descriptor in memory format	
FE60h	4 Bytes	GDTR	Reserved	Read-only
FE64h	Word		Limit	
FE66h	2 Bytes		Reserved	
FE68h	Quadword		Descriptor in memory format	
FE70h	Word	LDTR	Selector	Read-only
FE72h	Word		Attributes	
FE74h	Doubleword		Limit	
FE78h	Quadword		Base	

Table 6: SMM Save State

Offset	Size	Contents	Access
FE80h	4 Bytes	IDTR	Read-only
FE84h	Word	Reserved	
FEB6h	2 Bytes	Limit	
FE88h	Quadword	Reserved	
FE90h	Word	TR	Read-only
FE92h	Word	Selector	
FE94h	Doubleword	Attributes	
FE98h	Quadword	Limit	
FEA0h	Quadword	Base	Read-only
FEA8h	Quadword	IO_RESTART_RIP	
FEB0h	Quadword	IO_RESTART_RCX	
FEB8h	Quadword	IO_RESTART_RSI	
FEC0h	Doubleword	IO_RESTART_RDI	Read-only
FEC4	Doubleword	SMMFEC0 [SMM IO Trap Offset]	Read-only
FEC8h	Byte	SMMFEC4 [Local SMI Status]	Read-write
FEC9h	Byte	SMMFEC8 [SMM IO Restart Byte]	Read-write
FECAh	Byte	SMMFEC9 [Auto Halt Restart Offset]	Read-write
FECBh	5 Bytes	SMMFECA [NMI Mask]	Read-write
FED0h	Quadword	Reserved	
FED8h	Quadword	EFER	Read-only
FEE0h	Quadword	SMMFED8 [SMM SVM State]	Read-only
FEE8h	Quadword	Guest VMCB physical address	Read-only
FEF0h	16 Bytes	SVM Virtual Interrupt Control	Read-only
FEFCh	Doubleword	Reserved	
FF00h	Doubleword	SMMFEFC [SMM-Revision Identifier]	Read-only
FF04h	28 Bytes	SMMFF00 [SMM Base Address (SMM_BASE)]	Read-write
FF20h	Quadword	Reserved	
FF28h	Quadword	Guest PAT	Read-only
FF30h	Quadword	Host EFER ²	
FF38h	Quadword	Host CR4 ²	
FF40h	Quadword	Nested CR3 ²	
FF48h	Quadword	Host Cr0 ²	
FF50h	Quadword	CR4	
FF58h	Quadword	CR3	
FF60h	Quadword	CR0	
FF68h	Quadword	DR7	Read-only
FF70h	Quadword	DR6	
		RFLAGS	Read-write

Table 6: SMM Save State

Offset	Size	Contents	Access
FF78h	Quadword	RIP	Read-write
FF80h	Quadword	R15	
FF88h	Quadword	R14	
FF90h	Quadword	R13	
FF98h	Quadword	R12	
FFA0h	Quadword	R11	
FFA8h	Quadword	R10	
FFB0h	Quadword	R9	
FFB8h	Quadword	R8	
FFC0h	Quadword	RDI	
FFC8h	Quadword	RSI	
FFD0h	Quadword	RBP	
FFD8h	Quadword	RSP	
FFE0h	Quadword	RBX	
FFE8h	Quadword	RDX	
FFF0h	Quadword	RCX	
FFF8h	Quadword	RAX	

Notes:

1. This notation specifies that bit[47] is replicated in each of the 16 MSB's of the DW (sometimes called *sign extended*). The 16 LSB's contain bits[47:32].
2. Only used for an SMI in guest mode with nested paging enabled.

The SMI save state includes most of the integer execution unit. Not included in the save state are: the floating point state, MSRs, and CR2. In order to be used by the SMI handler, these must be saved and restored. The save state is the same, regardless of the operating mode (32-bit or 64-bit).

The following are some offsets in the SMM save state area. The mnemonic for each offset is in the form SMMxxxx, where xxxx is the offset in the save state.

SMMFEC0 SMM IO Trap Offset

If the assertion of SMI is recognized on the boundary of an IO instruction, [SMMFEC0 \[SMM IO Trap Offset\]](#) contains information about that IO instruction. For example, if an IO access targets an unavailable device, the system can assert SMI and trap the IO instruction. [SMMFEC0](#) then provides the SMI handler with information about the IO instruction that caused the trap. After the SMI handler takes the appropriate action, it can reconstruct and then re-execute the IO instruction from SMM. Or, more likely, it can use [SMMFEC8 \[SMM IO Restart Byte\]](#), to cause the core to re-execute the IO instruction immediately after resuming from SMM.

Bits	Description
31:16	Port: trapped IO port address. Read-only. This provides the address of the IO instruction.
15:12	BPR: IO breakpoint match. Read-only.
11	TF: EFLAGS TF value. Read-only.

10:7	Reserved
6	SZ32: size 32 bits. Read-only. 1=Port access was 32 bits.
5	SZ16: size 16 bits. Read-only. 1= Port access was 16 bits.
4	SZ8: size 8 bits. Read-only. 1=Port access was 8 bits.
3	REP: repeated port access. Read-only.
2	STR: string-based port access. Read-only.
1	V: IO trap word valid. Read-only. 1=The core entered SMM on an IO instruction boundary; all information in this offset is valid. 0=The other fields of this offset are not valid.
0	RW: port access type. Read-only. 0=IO write (OUT instruction). 1=IO read (IN instruction).

SMMFEC4 Local SMI Status

This offset stores status bits associated with SMI sources local to the core. For each of these bits, 1=The associated mechanism generated an SMI.

Bits	Description
31:20	Reserved.
19	SmiSrcThrCntHt: SMI source link thresholding. Read-only. This bit is associated with the SMI source specified in the link thresholding register (see MSR0000_0403 [LS Machine Check Miscellaneous (MC0_MISC)]).
18	Reserved.
17	SmiSrcLvtExt: SMI source LVT extended entry. Read-only. This bit is associated with the SMI sources specified in APIC[530:500] [Extended Interrupt [3:0] Local Vector Table] .
16	SmiSrcLvtLcy: SMI source LVT legacy entry. Read-only. This bit is associated with the SMI sources specified by the non-extended LVT entries of the APIC.
15:11	Reserved.
10	IntPendSmiSts: interrupt pending SMI status. Read-only. This bit is associated with the SMI source specified in MSRC001_0055 [Interrupt Pending][IntrPndMsg] (when that bit is high).
9	Reserved.
8	MceRedirSts: machine check exception redirection status. Read-only. This bit is associated with the SMI source specified in MSRC001_0022 [Machine Check Exception Redirection][RedirSmiEn] .
7:4	Reserved.
3:0	IoTrapSts: IO trap status. Read-only. Each of these bits is associated with each of the respective SMI sources specified in MSRC001_00[53:50] [IO Trap (SMI_ON_IO_TRAP_[3:0])] .

SMMFEC8 SMM IO Restart Byte

00h on entry into SMM.

If the core entered SMM on an IO instruction boundary, the SMI handler may write this to FFh. This causes the core to re-execute the trapped IO instruction immediately after resuming from SMM. The SMI handler should only write to this byte if [SMMFEC0\[V\]=1](#); otherwise, the behavior is undefined.

If a second SMI is asserted while a valid IO instruction is trapped by the first SMI handler, the core services the second SMI prior to re-executing the trapped IO instruction. [SMMFEC0\[V\]=0](#) during the second entry into

SMM, and the second SMI handler must not rewrite this byte.

If there is a simultaneous SMI IO instruction trap and debug breakpoint trap, the processor first responds to the SMI and postpones recognizing the debug exception until after resuming from SMM. If debug registers other than DR6 and DR7 are used while in SMM, they must be saved and restored by the SMI handler. If [SMMFEC8 \[SMM IO Restart Byte\]](#), is set to FFh when the RSM instruction is executed, the debug trap does not occur until after the IO instruction is re-executed.

Bits	Description
7:0	RST: SMM IO Restart Byte. Read-write.

SMMFEC9 Auto Halt Restart Offset

Bits	Description
7:1	Reserved.
0	HLT: halt restart. Read-write. Upon SMM entry, this bit indicates whether SMM was entered from the Halt state. 0=Entered SMM on a normal x86 instruction boundary. 1=Entered SMM from the Halt state. Before returning from SMM, this bit can be written by the SMI handler to specify whether the return from SMM should take the processor back to the Halt state or to the instruction-execution state specified by the SMM state save area (normally, the instruction after the halt). 0=Return to the instruction specified in the SMM save state. 1=Return to the halt state. If the return from SMM takes the processor back to the Halt state, the HLT instruction is not refetched and re-executed. However, the Halt special bus cycle is broadcast and the processor enters the Halt state.

SMMFECA NMI Mask

Bits	Description
7:1	Reserved.
0	NmiMask. Read-write. Specifies whether NMI was masked upon entry to SMM. 0=NMI not masked. 1=NMI masked.

SMMFED8 SMM SVM State

Read-only. This offset stores the SVM state of the processor upon entry into SMM.

Bits	Description
63:4	Reserved.

3	HostEflagsIf: host Eflags IF.	
2:0	SvmState.	
	<u>Bits</u>	<u>Definition</u>
	000b	SMM entered from a non-guest state.
	001b	Reserved.
	010b	SMM entered from a guest state.
	101b-011b	Reserved.
	110b	SMM entered from a guest state with nested paging enabled.
	111b	Reserved.

SMMFEFC SMM-Revision Identifier

SMM entry state: 0003_0064h

Bits	Description
31:18	Reserved.
17	BRL. Read-only. Base relocation supported.
16	IOTrap. Read-only. IO trap supported.
15:0	Revision. Read-only.

SMMFF00 SMM Base Address (SMM_BASE)

Bits	Description
31:0	See: MSRC001_0111 [SmmBase].

2.4.8.2.6 Exceptions and Interrupts in SMM

When SMM is entered, the core masks INTR, NMI, SMI, INIT, and A20M interrupts. The core clears the IF flag to disable INTR interrupts. To enable INTR interrupts within SMM, the SMM handler must set the IF flag to 1. A20M is disabled so that address bit 20 is never masked when in SMM.

Generating an INTR interrupt can be used for unmasking NMI interrupts in SMM. The core recognizes the assertion of NMI within SMM immediately after the completion of an IRET instruction. Once NMI is recognized within SMM, NMI recognition remains enabled until SMM is exited, at which point NMI masking is restored to the state it was in before entering SMM.

While in SMM, the core responds to the DBREQ and STPCLK interrupts, as well as to all exceptions that may be caused by the SMI handler.

2.4.8.2.7 The Protected ASeg and TSeg Areas

These ranges are controlled by [MSRC001_0112](#) and [MSRC001_0113](#); see those registers for details.

2.4.8.2.8 SMM Special Cycles

Special cycles can be initiated on entry and exit from SMM to acknowledge to the system that these transitions are occurring. These are controlled by [MSRC001_0015](#)[SMISPCYCDIS, RSMSPCYCDIS].

2.4.8.2.9 Locking SMM

The SMM registers ([MSRC001_0112](#) and [MSRC001_0113](#)) can be locked from being altered by setting [MSRC001_0015\[SmmLock\]](#). The BIOS can lock the SMM registers after initialization to prevent unexpected changes to these registers.

2.4.8.2.10 Synchronizing SMM Entry (Spring-Boarding)

The BIOS must take special care to ensure that all cores have entered SMM prior to accessing shared IO resources and all core SMI interrupt status bits are synchronized. This generally requires that BIOS waits for all cores to enter SMM.

The following conditions can cause one or more cores to enter SMM without all cores entering SMM:

- More than one IO device in the system is enabled to signal an SMI without hardware synchronization (e.g. using an end of SMI gate).
- A single device may signal multiple SMI messages without hardware synchronization (e.g. using an end of SMI gate).
- An SMI is received while one or more AP cores are in the INIT state. This may occur either during BIOS or secure boot.
- A hardware error prevents a core from entering SMM.

The act of synchronizing cores into SMM is called spring-boarding. Because not all of the above conditions can be avoided, it is recommended that all systems support spring-boarding.

An ACPI-compliant IO hub is required for spring-boarding. Depending on the IO hub design, BIOS may have to set additional end-of-SMI bits to trigger an SMI from within SMM.

The software requirements for the suggested spring-boarding implementation are listed as follows.

- A binary semaphore located in SMRAM, accessible by all cores. For the purpose of this discussion, the semaphore is called CheckSpringBoard. CheckSpringBoard is initialized to zero.
- Two semaphores located in SMRAM, accessible by all cores. For the purpose of this discussion, the semaphores are called NotInSMM and WaitInSMM. NotInSMM and WaitInSMM are initialized to a value equal to the number of cores in the system (NumCPUs).

The following BIOS algorithm describes spring-boarding and is optimized to reduce unnecessary SMI activity. This algorithm must be made part of the SMM instruction sequence for each core in the system.

1. Attempt to obtain ownership of the CheckSpringBoard semaphore with a read-modify-write instruction. If ownership was obtained then do the following, else proceed to step 2:

- Check all enabled SMI status bits in the IO hub.
Let Status=enable1&status1 | enable2&status2 | enable3&status3 ... enable n & status n.
- If (Status==0) then perform the following sub-actions.
 - Trigger an SMI broadcast assertion from the IO hub by writing to the software SMI command port.
 - Resume from SMM with the RSM instruction.

//Example:

```
InLineASM{
    BTS CheckSpringBoard,0; Try to obtain ownership of semaphore
    JC Step_2:
    CALL CheckIOHUB_SMIEVT; proc returns ZF=1 for no events
    JNZ Step_2:
    CALL Do_SpringBoard;Trigger SMI and then RSM
```

Step_2:

}

2. Decrement the NotInSMM variable. Wait for (NotInSMM==0). See Note 1.
3. Execute the core-local event SMI handler. Using a third semaphore (not described here), synchronize core execution at the end of the task. After all cores have executed, proceed to step 4. The following is a brief description of the task for each core:
 - Check all enabled core-local SMI status bits in the core's private or MSR address space. Handle the event if possible, or pass information necessary to handle the event to a mailbox for the BSC to handle.
 - An exclusive mailbox must exist for each core for each core local event.
 - On-line spare events should be handled in this task by the individual core for optimal performance. Assign one core of a dual core processor to handle On-line spare. These events may be optionally handled by the BSC just as other global events.
 - Wait for all cores to complete this task at least once.
4. If the current core executing instructions is not the BSC then jump to step 5. If the core executing instructions is the BSC then jump to the modified main SMI handler task, described below.
 - Check all enabled SMI status bits in the IO hub. Check mailboxes for event status.
 - For each event, handle the event and clear the corresponding status bit.
 - Repeat until all enabled SMI status bits are clear and no mailbox events remain.
 - Set NotInSMM=NumCPUs. (Jump to step 5.)
5. Decrement the WaitInSMM variable. Wait for WaitInSMM=0. See Note 2.
6. Increment the WaitInSMM variable. Wait for WaitInSMM=NumCPUs.
7. If the current processor core executing instructions is the BSC then reset CheckSpringBoard to zero.
8. Resume from SMM with the RSM instruction.

Notes:

1. To support a secure startup by the secure loader the BIOS must provide a timeout escape from the otherwise endless loop. The timeout value should be large enough to account for the latency of all cores entering SMM. The maximum SMM entrance latency is defined by the platform's IO sub-system, not the processor. A value of twice the watchdog timer count is recommended. See [D18F3x44 \[MCA NB Configuration\]](#) for more information on the watchdog time-out value.
If a time-out occurs in the wait loop, the BIOS (the last core to decrement NotInSMM) should record the number of cores that have not entered SMM and all cores must fall out of the loop.
2. If a time-out occurs in the wait loop in step 2, the BIOS must not wait for WaitInSMM=0. Instead it must wait for WaitInSMM="the number of cores recorded in step 2".

2.4.9 Secure Virtual Machine Mode (SVM)

Support for SVM mode is indicated by [CPUID Fn8000_0001_ECX\[SVM\]](#).

2.4.9.1 BIOS support for SVM Disable

The BIOS should include the following user setup options to enable and disable AMD Virtualization™ technology.

- Enable AMD Virtualization™.
 - [MSRC001_0114\[Svm_Disable\]](#) = 0.
 - [MSRC001_0114\[Lock\]](#) = 1.
 - [MSRC001_0118\[SvmLockKey\]](#) = 0000_0000_0000_0000h.
- Disable AMD Virtualization™.
 - [MSRC001_0114\[Svm_Disable\]](#)=1.
 - [MSRC001_0114\[Lock\]](#)=1.

- [MSRC001_0118](#)[SvmLockKey] = 0000_0000_0000_0000h.

The BIOS may also include the following user setup options to disable AMD Virtualization™.

- Disable AMD Virtualization™, with a user supplied key.
 - [MSRC001_0114](#)[Svm_Disable]=1.
 - [MSRC001_0114](#)[Lock]=1.
 - [MSRC001_0118](#)[SvmLockKey] programmed with value supplied by user. This value should be stored in NVRAM.

2.4.10 CPUID Instruction

The CPUID instruction provides data about the features supported by the processor. See [3.17 \[CPUID Instruction Registers\]](#).

2.4.10.1 Multi-Core Support

There are two methods for determining multi-core support. A recommended mechanism is provided and a legacy method is also available for existing operating systems. System software should use the correct architectural mechanism to detect the number of physical cores by observing [CPUID Fn8000_0008_ECX](#)[NC]. The legacy method utilizes the [CPUID Fn0000_0001_EBX](#)[LogicalProcessorCount].

2.5 Power Management

The processor supports many power management features in a variety of systems. [Table 7](#) provides a summary of ACPI states and power management features and indicates whether they are supported.

Table 7: Power Management Support

ACPI/Power Management State	Supported	Description
G0/S0/C0: Working	Yes	
G0/S0/C0: Core P-state transitions	Yes	2.5.3.1 [Core P-states]
G0/S0/C0: NB P-state transitions	Yes	2.5.4.1 [NB P-states]
G0/S0/C0: Hardware thermal control (HTC)	Yes	2.10.4.1 [PROCHOT_L and Hardware Thermal Control (HTC)]
G0/S0/C0: P-state limit control	Yes	2.10.4.3 [Software P-state Limit Control]
G0/S0/C0: Thermal clock throttling (SMC controlled)	No	
G0/S0/Per-core IO-based C-states	Yes	2.5.3.2 [Core C-states]
G0/S0/C1: Halt	Yes	
G0/S0/PC4: Altvid (VDD power plane)	No	
G0/S0/C5: Deeper altvid support (VDD power plane)	No	
G0/S0/CC6: Power gating	Yes	
G0/S0/Cx: Cache flushing support	Yes	
G1/S1: Stand By (Powered On Suspend)	No	
G1/S3: Stand By (Suspend to RAM)	Yes	2.5.8.1 [S-states]
G1/S4, S5: Hibernate (Suspend to Disk), Shut Down (Soft Off)	Yes	
G3 Mechanical Off	Yes	
Parallel VID Interface	No	2.5.1 [Processor Power Planes And Voltage Control]
Serial VID Interface 1	No	
Serial VID Interface 2	Yes	
Single-plane systems	No	
Dual-plane systems	Yes	
Triple-plane systems	No	

2.5.1 Processor Power Planes And Voltage Control

Refer to the *Electrical Data Sheet for AMD Family 15h Models 10h-1Fh Processors* for power plane definitions.

2.5.1.1 Serial VID Interface

The processor includes an interface to control external voltage regulators, called the serial VID interface (SVI). Only SVI2 is supported. The frequency of SVC for SVI2 is controlled by [D18F3xA0\[Svi2HighFreqSel\]](#).

2.5.1.1.1 SVI2 Features

The processor supports the following SVI2 features:

- Voltage offsets:
 - VDD: [D18F5x12C\[CoreOffsetTrim\]](#).
 - VDDNB: [D18F5x188\[NbOffsetTrim\]](#).

- Load line trim:
 - VDD: [D18F5x12C](#)[CoreLoadLineTrim].
 - VDDNB: [D18F5x188](#)[NbLoadLineTrim].

2.5.1.2 Internal VID Registers and Encodings

All VID register fields within the processor are 8-bits wide. The SVI2 VID encoding to voltage level is calculated by the following formula:

```
IF (Vid[7:0] < 1111_1000b) THEN
  Voltage (V) = 1.5500 - 0.00625 * Vid[7:0]
ELSE
  Voltage (V) = OFF
END
```

The boot VID is 1.0 volts.

2.5.1.2.1 MinVid and MaxVid Check

Hardware limits the minimum and maximum VID code that is sent to the voltage regulator. The allowed limits are specified in [D18F5x17C](#)[MinVid, MaxVid]. Prior to generating VID-change commands to SVI, the processor filters the InputVid value to the OutputVid as follows (higher VID codes correspond to lower voltages and lower VID codes correspond to higher voltages):

- If InputVid < MaxVid, OutputVid=MaxVid.
 - Else if (InputVid > MinVid) & (MinVid!=00h), OutputVid=MinVid.
 - Else OutputVid=InputVid.

This filtering is applied regardless of the source of the VID-change command.

2.5.1.3 Low Power Features

2.5.1.3.1 PSIx_L Bit

The processor supports indication of whether the processor is in a low-voltage state or not, which may be used by the regulator to place itself into a more power efficient mode. The PSIx_L bit can be controlled for the VDD and VDDNB power planes independently.

- The processor supports the PSIO_L and the PSI1_L bits in the data fields of the SVI2 command.
 - PSIO_L: PSIO_L is enabled using [D18F3xA0](#)[PsiVidEn] and [D18F5x17C](#)[NbPsi0VidEn]. Once enabled, the state of PSIO_L is controlled by [D18F3xA0](#)[PsiVid[7:0]] and [D18F5x17C](#)[NbPsi0Vid[7:0]].
 - PSI1_L: The PSI1_L bit for VDD functions as specified by [D18F5x12C](#)[CorePsi1En]. The PSI1_L bit for VDDNB functions as specified by [D0F0xBC_x1F5F8](#)[EnableNbPsi1].

The processor sends PSIx_L changes to the voltage regulator whenever the criteria associated with PSIx_L changes, no voltage change is necessary.

2.5.1.3.1.1 BIOS Requirements for PSIO_L

Enabling PSIO_L for the VDD and VDDNB planes depends on support from the voltage regulator and is therefore system specific. The voltage regulator must be able to supply the current required for the processor to operate at the VID code specified in [D18F3xA0](#)[PsiVid[7:0]] and [D18F5x17C](#)[NbPsi0Vid[7:0]]. Depending on the regulator used, AMD recommends one of the following methods:

- PSIO_L is disabled:
 - VDD: To set PSIO_L for the VDD plane, program `D18F3xA0[PsiVidEn]=0`.
 - VDDNB: To set PSIO_L for the VDDNB plane, program `D18F5x17C[NbPsi0VidEn]=0`.
- PSIO_L set/clear based on current requirements:
 - VDD: The following algorithm describes how to program PSIO_L on VDD:

```

PSI_vrm_current = current at which the regulator allows PSIO_L.
previous_voltage = FFh

for (each P-state from P0 to D18F3xDC[HwPstateMaxVal]) {
  pstate_current = ProcIddMax for the current P-state,
                  see 2.5.3.1.8 [Processor-Systemboard Power Delivery Compatibility Check];
  pstate_voltage = MSRC001_00[6B:64][CpuVid] of the current P-state;

  if (current P-state == D18F3xDC[HwPstateMaxVal]) {
    next_pstate_current = 0;
  } else {
    next_pstate_current = ProcIddMax for the next P-state,
                        see 2.5.3.1.8 [Processor-Systemboard Power Delivery Compatibility Check];
  }

  if ((pstate_current <= PSI_vrm_current) &&
      (next_pstate_current <= PSI_vrm_current) &&
      (pstate_voltage != previous_voltage)) {
    Program D18F3xA0[PsiVid] = pstate_voltage;
    Program D18F3xA0[PsiVidEn] = 1;
    break;
  }
  previous_voltage = pstate_voltage;
}

```

- VDDNB: The following algorithm describes how to program PSIO_L on VDDNB:

```

NbIddMax = D18F5x1[6C:60][NbIddDiv] current.

PSI_vrm_current = current at which the VDDNB regulator allows PSIO_L.
previous_voltage = FFh

for (each valid NB P-state starting with NBP0) {
  pstate_current = NbIddMax of the current NB P-state;
  pstate_voltage = D18F5x1[6C:60][NbVid] of the current NB P-state;

  if (current NB P-state is the last valid NB P-state) {
    next_pstate_current = 0;
  } else {
    next_pstate_current = NbIddMax for the next NB P-state;
  }

  if ((pstate_current <= PSI_vrm_current) &&
      (next_pstate_current <= PSI_vrm_current) &&
      (pstate_voltage != previous_voltage)) {
    Program D18F5x17C[NbPsi0Vid] = pstate_voltage;
    Program D18F5x17C[NbPsi0VidEn] = 1;
    break;
  }
  previous_voltage = pstate_voltage;
}

```

2.5.1.3.1.2 BIOS Requirements for PS11_L

- VDD: Program [D18F5x12C](#)[CorePsi1En] to its recommended setting.
- VDDNB: Please see your AMD representative for details.

2.5.1.3.2 Low Power Voltages

In order to save power, voltages lower than those normally needed for operation may be applied to the VDD power plane while the processor is in a C-state or S-state. The lower voltage is defined as follows:

- PC6Vid: [D18F5x128](#)[PC6Vid] specifies a voltage that does not retain the CPU caches or the core microarchitectural state. PC6Vid does not allow execution and is only applied to the cores. See [2.5.3.2.3.4 \[Package C6 \(PC6\) State\]](#).

2.5.1.4 Voltage Transitions

The processor supports dynamic voltage transitions on the VDD and VDDNB planes. These transitions are requested by either hardware or software during state changes such as reset, P-state changes, and C-state changes. In all cases the VID code passed to the voltage regulator changes from the old value to the new value without stepping through intermediate values. The voltage regulator ramps the voltage directly from the starting voltage to the final voltage, no stepping occurs. See the *AMD Serial VID Interface 2.0 (SVI2) Specification* for additional details.

- If a voltage increase is requested, the processor waits as specified by [D18F5x12C](#)[WaitVidCompDis] before sending any additional voltage change requests to the voltage regulator or before beginning a frequency transition.
- If a voltage decrease is requested, the processor waits the amount of time specified by [D18F5x128](#)[FastSlamTimeDown] before sending any additional voltage change requests to the voltage regulator. For voltage decreases, the processor does not wait any time before beginning frequency changes.

The processor continues code execution during voltage changes when in the C0 state.

2.5.1.4.1 Hardware-Initiated Voltage Transitions

When software requests any of the following state changes, or hardware determines that any of the following state changes are necessary, hardware coordinates the necessary voltage changes:

- VDD:
 - Core P-state transition. See [2.5.3.1 \[Core P-states\]](#).
 - Package C-state transition. See [2.5.3.2 \[Core C-states\]](#).
 - S-state transition. See [2.5.8.1 \[S-states\]](#).
- VDDNB:
 - NB P-state transition. See [2.5.4.1 \[NB P-states\]](#).
 - S-state transition. See [2.5.8.1 \[S-states\]](#).

2.5.1.4.2 Software-Initiated Voltage Transitions

2.5.1.4.2.1 Software-Initiated NB Voltage Transitions

Software can request voltage changes on the VDDNB power plane using the following control/status register

pairs:

- [GMMx63C](#) and [GMMx640](#)
- [GMMx770](#) and [GMMx774](#)

The voltage requests from each register pair are considered independently by hardware when taking voltage plane dependencies into account (see [2.5.2.2 \[Dependencies Between Subcomponents on VDDNB\]](#)). To make a voltage change request, software uses the following sequence:

1. Ensure `VoltageChangeEn==1` in the control register. If software needs to program `VoltageChangeEn=1`, software must perform this register write independently of the writes in the following steps.
2. Program `VoltageLevel` to the desired voltage and toggle `VoltageChangeReq` in the control register.
3. The voltage change is complete when `VoltageChangeReq` in the control register is equal to `VoltageChangeAck` in the status register.

Software can force a VDDNB voltage change using [GMMx63C](#) and [GMMx770](#). To do so, software programs `VoltageForceEn=1` in the respective register before toggling `VoltageChangeReq` when making a voltage change request. If this is done, the voltage requested overrides any other VDDNB voltage requests made by software or DPM state transitions when determining voltage plane dependencies (see [2.5.2.2 \[Dependencies Between Subcomponents on VDDNB\]](#)). NB P-state transitions still request voltage transitions as normal when software forces a voltage change using this mechanism. If software forces a voltage change using both [GMMx63C](#) and [GMMx770](#), the voltage requested in [GMMx63C](#) takes precedence.

The following registers also cause VDDNB voltage transitions:

- [MSRC001_0070\[NbVid\]](#): See [2.5.4.1 \[NB P-states\]](#).

2.5.1.4.2.2 Software-Initiated Core Voltage Transitions

To force VDD voltage changes, software must take the following steps:

1. Write the destination `CpuVid` to [MSRC001_0070\[CpuVid\]](#).
2. Wait the specified [D18F3xD8\[VSRampSlamTime\]](#).

2.5.2 Frequency and Voltage Domain Dependencies

2.5.2.1 Dependencies Between Cores

Whenever a P-state or C-state is requested on a core (see [2.5.3.1 \[Core P-states\]](#) and [2.5.3.2 \[Core C-states\]](#)), hardware must take the following frequency and voltage domain dependencies into account when deciding whether to make the requested change:

- Cores within a compute unit share a common frequency and voltage domain.
- Compute units within a processor share a common voltage domain, but have independent frequency domains. The voltage is determined by the highest-performance P-state requested on any core.

As a result, the P-state and C-state change requests have the following results:

- If different compute units request different voltages, the VDD voltage is determined by the highest voltage (lowest VID) requested.
- If the cores within a compute unit request different P-states while in C0, frequency and voltage are deter-

mined by the highest-performance P-state requested.

- If one core within a compute unit requests a C-state while the other core is in C0, the frequency and voltage of the compute unit is determined by the core in C0.
- If both cores request non-C0 states, the behavior is specified by [D18F4x128\[CoreCstatePolicy\]](#).

2.5.2.2 Dependencies Between Subcomponents on VDDNB

Many subcomponents of the processor including the NB, the GPU, and the root complex reside on the VDDNB power plane. Hardware must take voltage domain dependencies into account when determining whether to make a voltage change requested by one of the subcomponents. Whenever a state transition occurs that causes a voltage change request (see [2.5.1.4.1 \[Hardware-Initiated Voltage Transitions\]](#)), or software makes a voltage change request (see [2.5.1.4.2 \[Software-Initiated Voltage Transitions\]](#)), the VDDNB voltage requested by the processor is determined by the highest voltage (lowest VID) request made by any of the subcomponents or by software. In addition, software can force VDDNB voltage changes. See [2.5.1.4.2 \[Software-Initiated Voltage Transitions\]](#).

2.5.2.3 BIOS Requirements for Power Plane Initialization

- Ensure the following fields are configured to their BIOS recommendations:
 - [D18F3xA0\[Svi2HighFreqSel\]](#).
 - [D18F3xD8\[VSRampSlamTime\]](#).

2.5.3 CPU Power Management

2.5.3.1 Core P-states

Core P-states are operational performance states characterized by a unique combination of core frequency and voltage. The processor supports up to 8 core P-states (P0 through P7), specified in [MSRC001_00\[6B:64\]](#). Out of cold reset, the voltage and frequency of the compute units is specified by [MSRC001_0071\[StartupPstate\]](#).

Support for dynamic core P-state changes is indicated by more than one enabled selection in [MSRC001_00\[6B:64\]\[PstateEn\]](#). At least one enabled P-state (P0) is specified for all processors.

Software requests core P-state changes for each core independently using the hardware P-state control mechanism (a.k.a. fire and forget). Support for hardware P-state control is indicated by [CPUID Fn8000_0007_EDX\[HwPstate\]=1b](#). Software may not request any P-state transitions using the hardware P-state control mechanism until the P-state initialization requirements defined in [2.5.3.1.7 \[BIOS Requirements for Core P-state Initialization and Transitions\]](#) are complete.

2.5.3.1.1 Application Power Management (APM)

Application Power Management (APM) allows the processor to deterministically provide maximum performance while remaining within the specified power delivery and removal envelope. APM dynamically monitors processor activity and generates an approximation of power consumption. If power consumption exceeds a defined power limit, a P-state limit is applied by APM hardware to reduce power consumption. APM ensures that average power consumption over a thermally significant time period remains at or below the defined power limit. This allows P-states to be defined with higher frequencies and voltages than could be used without APM. These P-states are referred to as boosted P-states.

- Support for APM is specified by [CPUID Fn8000_0007_EDX\[CPB\]](#).
- APM is enabled if all of the following conditions are true:
 - [MSRC001_0015\[CpbDis\]](#) = 0 for all cores.

- `D18F4x15C[ApmMasterEn]` = 1.
- `D18F4x15C[BoostSrc]` = 01b.
- `D18F4x15C[NumBoostStates]` != 0.
- APM can be dynamically enabled and disabled through `MSRC001_0015[CpbDis]`. If core performance boost (CPB) is disabled, a P-state limit is applied. The P-state limit restricts cores to the highest performance non-boosted P-state.
- All P-states, both boosted and non-boosted, are specified in `MSRC001_00[6B:64]`.
- The number of boosted P-states is specified by `D18F4x15C[NumBoostStates]`.
 - The number of boosted P-states may vary from product to product.
- There are two stages of boost supported. Compute units can be placed in the first stage of boosted P-states if the processor power consumption remains within the TDP limit. The second stage of boosted P-states can only be achieved if a subset of compute units are in CC6 and the processor power consumption remains within the TDP limit. See `D18F4x16C[CstateCnt, CstateBoost]`.
- All boosted P-states are always higher performance than non-boosted P-states.
- To ensure proper operation, boosted P-states should be hidden from the operating system. BIOS should not provide ACPI `_PSS` entries for boosted P-states. See 2.5.3.1.9.3.2 [`_PSS (Performance Supported States)`].
- The lowest-performance P-state CPB limits the processor to is the highest-performance non-boosted P-state.

2.5.3.1.2 Core P-state Naming and Numbering

Since the number of boosted P-states may vary from product to product, the mapping between `MSRC001_00[6B:64]` and the indices used to request P-state changes or status also varies. In order to clarify this, two different numbering schemes are used.

2.5.3.1.2.1 Software P-state Numbering

When referring to software P-state numbering, the following naming convention is used:

- Non-boosted P-states are referred to as P0, P1, etc.
 - P0 is the highest power, highest performance, non-boosted P-state.
 - Each ascending P-state number represents a lower-power, lower performance non-boosted P-state than the prior P-state number.
- Boosted P-states are referred to as Pb0, Pb1, etc.
 - Pb0 is the highest-performance, highest-power boosted P-state.
 - Each higher numbered boosted P-state represents a lower-power, lower-performance boosted P-state.

For example, if `D18F4x15C[NumBoostStates]` contains the values shown below, then the P-states would be named as follows:

Table 8: Software P-state Naming

<code>D18F4x15C[NumBoostStates]=1</code>		<code>D18F4x15C[NumBoostStates]=3</code>	
P-state Name	Corresponding MSR Address	P-state Name	Corresponding MSR Address
Pb0	MSRC001_0064	Pb0	MSRC001_0064
P0	MSRC001_0065	Pb1	MSRC001_0065
P1	MSRC001_0066	Pb2	MSRC001_0066
P2	MSRC001_0067	P0	MSRC001_0067
P3	MSRC001_0068	P1	MSRC001_0068

Table 8: Software P-state Naming

D18F4x15C[NumBoost-States]=1		D18F4x15C[NumBoost-States]=3	
P-state Name	Corresponding MSR Address	P-state Name	Corresponding MSR Address
P4	MSRC001_0069	P2	MSRC001_0069
P5	MSRC001_006A	P3	MSRC001_006A
P6	MSRC001_006B	P4	MSRC001_006B

All sections and register definitions use software P-state numbering unless otherwise specified.

2.5.3.1.2.2 Hardware P-state Numbering

When referring to hardware P-state numbering, the following naming convention is used:

- All P-states are referred to as P0, P1, etc.
 - P0 is the highest power, highest-performance P-state, regardless of whether it is a boosted P-state or a non-boosted P-state.
 - Each ascending P-state number represents a lower-power, lower-performance P-state, regardless of whether it is a boosted P-state or not.

2.5.3.1.3 Core P-state Control

Core P-states are dynamically controlled by software and are exposed through ACPI objects (refer to [2.5.3.1.9.3 \[ACPI Processor P-state Objects\]](#)). Software requests a core P-state change by writing a 3 bit index corresponding to the desired P-state number to `MSRC001_0062[PstateCmd]` of the appropriate core. For example, to request P3 for core 0 software would write 011b to core 0's `MSRC001_0062[PstateCmd]`. Boosted P-states may not be directly requested by software. Whenever software requests the P0 state on a processor that supports CPB (i.e. writes 000b to `MSRC001_0062[PstateCmd]`), hardware dynamically places the core into the highest-performance P-state possible as determined by CPB. See [2.5.3.1.1 \[Application Power Management \(APM\)\]](#).

Hardware sequences the frequency and voltage changes necessary to complete a P-state transition as specified by [2.5.3.1.6 \[Core P-state Transition Behavior\]](#) with no additional software interaction required. Hardware also coordinates frequency and voltage changes when differing P-state requests are made on cores that share a frequency or voltage plane. See [2.5.2 \[Frequency and Voltage Domain Dependencies\]](#) for details about hardware coordination.

Table 9: Software P-state Control

D18F4x15C[NumBoostStates]=1			D18F4x15C[NumBoostStates]=3		
P-state Name	Index Used for Requests/Status	Corresponding MSR Address	P-state Name	Index Used for Requests/Status	Corresponding MSR Address
Pb0	n/a	MSRC001_0064	Pb0	n/a	MSRC001_0064
P0	0	MSRC001_0065	Pb1	n/a	MSRC001_0065
P1	1	MSRC001_0066	Pb2	n/a	MSRC001_0066
P2	2	MSRC001_0067	P0	0	MSRC001_0067
P3	3	MSRC001_0068	P1	1	MSRC001_0068

Table 9: Software P-state Control

D18F4x15C[NumBoostStates]=1			D18F4x15C[NumBoostStates]=3		
P-state Name	Index Used for Requests/Status	Corresponding MSR Address	P-state Name	Index Used for Requests/Status	Corresponding MSR Address
P4	4	MSRC001_0069	P2	2	MSRC001_0069
P5	5	MSRC001_006A	P3	3	MSRC001_006A
P6	6	MSRC001_006B	P4	4	MSRC001_006B

Hardware controls the VID for each voltage domain according to the highest requirement of the frequency domain(s) on each plane. For example, the VID for a 4 compute unit dual-plane system must be maintained at the highest level required for all 4 frequency domains. The number of frequency domains in a voltage domain is package/platform specific. Refer to 2.5.3.1.6 [Core P-state Transition Behavior] for details on hardware P-state voltage control. 2.5.2.3 [BIOS Requirements for Power Plane Initialization] specifies the processor initialization requirements for voltage plane control.

2.5.3.1.4 Core P-state Visibility

MSRC001_0063 [P-state Status][CurPstate] reflects the current non-boosted P-state number for each compute unit. For example, if MSRC001_0063[CurPstate]=010b on compute unit 1, then compute unit 1 is in the P2 state. If a compute unit is in a boosted P-state, MSRC001_0063[CurPstate] reads back as 0.

The voltage on a compute unit may not correspond to the VID code specified by the current P-state of the compute unit due to voltage plane dependencies. See 2.5.2 [Frequency and Voltage Domain Dependencies]. If a compute unit is in the P0 state (i.e. if MSRC001_0063[CurPstate]=0), the frequency of the compute unit could be the frequency specified by P0 or any boosted P-state. To determine the frequency of a compute unit, see 2.5.3.3 [Effective Frequency].

2.5.3.1.5 Core P-state Limits

Core P-states may be limited to lower-performance values under certain conditions, including:

- HTC. See D18F3x64 [Hardware Thermal Control (HTC)][HtcPstateLimit].
- Software. See D18F3x68 [SwPstateLimit].
- Core Performance Boost. See 2.5.3.1.1 [Application Power Management (APM)].
- PROCHOT_L assertion. See 2.10.4.1 [PROCHOT_L and Hardware Thermal Control (HTC)].
- SMU. See D18F4x13C [SmuPstateLimit].

P-state limits are applied to all cores on the processor. The current P-state limit is provided in MSRC001_0061 [P-state Current Limit][CurPstateLimit]. Changes to the MSRC001_0061[CurPstateLimit] can be programmed to trigger interrupts through D18F3x64 [PslApicLoEn and PslApicHiEn]. In addition, the maximum P-state value, regardless of the source, is limited as specified in MSRC001_0061 [PstateMaxVal].

2.5.3.1.6 Core P-state Transition Behavior

The following rules specify how P-states changes function and interact with other system or processor states:

- If the P-state number is increasing (the compute unit is moving to a lower-performance state), then the COF is changed first, followed by the VID change. If the P-state number is decreasing, then the VID is changed first followed by the COF.
- When the processor initiates a VID change that increases voltage for a voltage domain, no new voltage or frequency changes occur until D18F3xD8 [VSRampSlamTime] has expired, regardless of whether any new

requests are received. When the processor initiates a VID change that decreases voltage for a voltage domain, new voltage or frequency changes are allowed to occur immediately.

- This is true regardless of whether the frequency or voltages changes occur as a result of P-state or C-state changes.
- If multiple commands are issued that affect the P-state of a domain prior to when the processor initiates the change of the P-state of that domain, then the processor operates on the last one issued.
- Once a P-state change starts, the P-state state machine (PSSM) continues through completion unless interrupted by a PWROK deassertion. If multiple P-state changes are requested concurrently, the PSSM may group the associated VID changes separately from the associated COF changes.
- Behavior during RESET_L assertions:
 - All compute units are transitioned to C0.
 - If there is no P-state transition activity, then the compute units and NB remain in the current P-state. If a RESET_L assertion interrupts a P-state transition, then the COF remains in it's current state at the time RESET_L is asserted (either the value of the old or the new P-state) and the VID remains in it's current state (perhaps at a VID between the old and the new P-states, if the VID was being stepped). BIOS is required to transition to valid COF and VID settings after a warm reset according to the sequence defined in 2.5.3.1.9 [BIOS COF and VID Requirements After Warm Reset].
 - After a warm reset MSRC001_0063 [P-state Status] is consistent with MSRC001_0071 [CurPstate]. MSRC001_0062 [P-state Control] may not be consistent with MSRC001_0071 [CurPstate].
 - If D18F5x1[6C:60][NbFid] has changed, then the new value is applied to the NB PLL on the assertion of RESET_L. It is assumed that BIOS adjusts the NB VID to the appropriate value prior to the warm reset. See 2.5.1.4.2 [Software-Initiated Voltage Transitions].
- The OS controls the P-state through MSRC001_0062 [P-state Control], independent of P-state limits described in D18F3x64 [Hardware Thermal Control (HTC)][HtcPstateLimit], D18F3x68 [Software P-state Limit][StcPstateLimit]. P-state limits interact with OS-directed P-state transitions as follows:
 - Of all the active P-state limits, the one that represents the lowest-performance P-state number, at any given time, is treated as an upper limit on performance.
 - As the limit becomes active or inactive, or if it changes, the P-state for each compute unit is placed in either the last OS-requested P-state or the new limit P-state, whichever is a lower performance P-state number.
 - If the resulting P-state number exceeds MSRC001_0061 [P-state Current Limit][PstateMaxVal], regardless of whether it is a limit or OS-requested, then the PstateMaxVal is used instead.

2.5.3.1.7 BIOS Requirements for Core P-state Initialization and Transitions

1. Check that CPUID Fn8000_0007_EDX[HwPstate]=1. If not, P-states are not supported and BIOS skips the rest of these steps.
2. Complete the 2.5.2.3 [BIOS Requirements for Power Plane Initialization].
3. Ensure the following fields are configured to their BIOS recommendations:
 - D18F3xA0[PIILockTime].
 - D18F3xD4[PowerStepUp, PowerStepDown].
4. Transition all cores to the minimum performance P-state using the algorithm detailed in 2.5.3.1.9.2 [Core Minimum P-state Transition Sequence After Warm Reset].
5. Complete the 2.5.4.1.2.1 [NB P-state COF and VID Synchronization After Warm Reset]. All cores on a processor must be in the minimum performance P-state prior to executing this sequence.
6. Complete the 2.5.3.1.8 [Processor-Systemboard Power Delivery Compatibility Check].
7. Perform the following steps in any order:
 - A. Enable 2.5.3.1.1 [Application Power Management (APM)] as follows:
 - Ensure the following fields are configured to their BIOS recommendations:
 - D18F4x110[CSampleTimer].
 - D18F4x15C[ApmMasterEn].

- [D18F5xE0\[RunAvgRange\]](#).
 - See your AMD representative for details on how to enable the GPU aspects of [2.5.3.1.1 \[Application Power Management \(APM\)\]](#).
 - If [D18F4x15C\[NumBoostStates\]!](#)=0, program [D18F4x15C\[BoostSrc\]=1](#).
- B. Transition all cores to the maximum performance P-state by writing 0 to [MSRC001_0062\[PstateCmd\]](#).
- C. Create ACPI objects if necessary:
- Determine the valid set of P-states as indicated by [MSRC001_00\[6B:64\]\[PstateEn\]](#).
 - If P-states are not supported, as indicated by only one enabled selection in [MSRC001_00\[6B:64\]\[PstateEn\]](#), then BIOS must not generate ACPI-defined P-state objects described in [2.5.3.1.9.3 \[ACPI Processor P-state Objects\]](#). Otherwise, the ACPI objects should be generated to enable P-state support.
- D. Configure the COF and VID for each processor appropriately based on the sequence described in [2.5.4.1.2 \[BIOS NB P-state Configuration\]](#).
8. Configure [PSIx_L](#). Refer to [2.5.1.3.1 \[PSIx_L Bit\]](#) for additional details.

2.5.3.1.8 Processor-Systemboard Power Delivery Compatibility Check

BIOS must disable processor P-states that require higher power delivery than the systemboard can support. This power delivery compatibility check is designed to prevent system failures caused by exceeding the power delivery capability of the systemboard for the power plane(s) that contain the core(s). Refer to [2.5.1 \[Processor Power Planes And Voltage Control\]](#) for power plane definitions and configuration information. BIOS can optionally notify the user if P-states are detected that exceed the systemboard power delivery capability. Modifications to [MSRC001_00\[6B:64\] \[P-state \[7:0\]\]](#) must be applied equally to all cores on the same node. This check does not ensure functionality for all package/socket compatible processor/systemboard combinations.

[MSRC001_00\[6B:64\]\[PstateEn\]](#) must be set to 0 for any P-state MSR where $PstateEn=1$ and the processor current requirement ($ProcIddMax$), defined by the following equation, is greater than the systemboard current delivery capability. $ProcIddMax = MSRC001_00[6B:64][IddValue] \text{ current} * 1/10^{MSRC001_00[6B:64][IddDiv]} * (D18F5x84[CmpCap]+1)$;

The power delivery check should be applied starting with hardware P0 and continue with increasing P-state indexes (1, 2, 3, and 4) for all enabled P-states. Once a compatible P-state is found using the $ProcIddMax$ equation the check is complete. All processor P-states with higher indexes are defined to be lower power and performance, and are therefore compatible with the systemboard.

Example:

- $MSRC001_0065[IddValue] = 32d$
- $MSRC001_0065[IddDiv] = 0d$
- $D18F5x84[CmpCap] = 1d$
- $ProcIddMax = 32 * 1 * 2 = 64A$ per plane

The systemboard must be able to supply $\geq 64A$ for the unified core power plane in order to support P1 for this processor. If the systemboard current delivery capability is $< 64A$ per plane then BIOS must set [MSRC001_0065\[PstateEn\]=0](#) for all cores on this node, and continue by checking P2 in the same fashion.

If no P-states are disabled while performing the power delivery compatibility check then BIOS does not need to take any action.

If at least one P-state is disabled by performing the power delivery compatibility check and at least one P-state remains enabled, then BIOS must perform the following steps:

1. If the P-state pointed to by [MSRC001_0063\[CurPstate\]](#) is disabled by the power delivery compatibility

- check, then BIOS must request a transition to an enabled P-state using `MSRC001_0062[PstateCmd]` and wait for `MSRC001_0063[CurPstate]` to reflect the new value.
- Copy the contents of the enabled P-state MSRs (`MSRC001_006B:64`) to the highest performance P-state locations. E.g. if P0 and P1 are disabled by the power delivery compatibility check and P2 - P4 remain enabled, then the contents of P2 - P4 should be copied to P0 - P2 and P3 and P4 should be disabled (`PstateEn=0`). This step uses software P-state numbering. See 2.5.3.1.2.1 [Software P-state Numbering].
 - Request a P-state transition to the P-state MSR containing the COF/VID values currently applied. E.g. If `MSRC001_0063[CurPstate]=100b` and P4 P-state MSR information is copied to P2 in step 2, then BIOS should write 010b to `MSRC001_0062[PstateCmd]` and wait for `MSRC001_0063[CurPstate]` to reflect the new value.
 - If a subset of boosted P-states are disabled, then copy the contents of the P-state MSR pointed to by the highest performance boosted P-state that is enabled to the P-state MSRs pointed to by the boosted P-states that are disabled.
 - If all boosted P-states are disabled, then program `D18F4x15C[BoostSrc]=0`.
 - Adjust the following P-state parameters affected by the P-state MSR copy by subtracting the number of software P-states that are disabled by the power delivery compatibility check. This calculation should not wrap, but saturate at 0. E.g. if P0 and P1 are disabled, then each of the following register fields should have 2 subtracted from them:
 - `D18F3x64[HtcPstateLimit]`
 - `D18F3x68[SwPstateLimit]`
 - `D18F3xDC[HwPstateMaxVal]`

If any node has all P-states disabled after performing the power delivery compatibility check, then BIOS must perform the following steps. This does not ensure operation and BIOS should notify the user of the incompatibility between the processor and systemboard if possible.

- If `MSRC001_0063[CurPstate] != MSRC001_0061[PstateMaxVal]`, then write `MSRC001_0061[PstateMaxVal]` to `MSRC001_0062[PstateCmd]` and wait for `MSRC001_0063[CurPstate]` to reflect the new value.
- If `MSRC001_0061[PstateMaxVal] != 000b` copy the contents of the P-state MSR pointed to by `MSRC001_0061[PstateMaxVal]` to `MSRC001_0064` and set `MSRC001_0064[PstateEn]`; Write 000b to `MSRC001_0062[PstateCmd]` and wait for `MSRC001_0063[CurPstate]` to reflect the new value. This step uses software P-state numbering. See 2.5.3.1.2.1 [Software P-state Numbering].
- Adjust the following fields to 000b.
 - `D18F3x64[HtcPstateLimit]`
 - `D18F3x68[SwPstateLimit]`
 - `D18F3xDC[HwPstateMaxVal]`
- Program `D18F4x15C[BoostSrc]=0`.

2.5.3.1.9 BIOS COF and VID Requirements After Warm Reset

Warm reset is asynchronous and can interrupt P-state transitions leaving the processor in a VID state that does not correspond to `MSRC001_0063[CurPstate]` on any core. The processor frequency after warm reset corresponds to `MSRC001_0063[CurPstate]`. See 2.5.3.1.6 [Core P-state Transition Behavior] for P-state transition behavior when `RESET_L` is asserted. BIOS is required to transition the processor to valid COF and VID settings corresponding to an enabled P-state following warm reset. The cores may be transitioned to either the maximum or minimum P-state COF and VID settings using the sequences defined in 2.5.3.1.9.1 [Core Maximum P-state Transition Sequence After Warm Reset] and 2.5.3.1.9.2 [Core Minimum P-state Transition Sequence After Warm Reset]. Transitioning to the minimum P-state after warm reset is recommended to prevent undesired system behavior if a warm reset occurs before the 2.5.3.1.8 [Processor-Systemboard Power Delivery Compatibility Check] is complete. BIOS is not required to manipulate NB COF and VID settings following warm reset if the warm reset was issued by BIOS to update `D18F5x1[6C:60][NbFid]`.

2.5.3.1.9.1 Core Maximum P-state Transition Sequence After Warm Reset

1. Write `MSRC001_0061[PstateMaxVal]` to `MSRC001_0062[PstateCmd]` on all cores in the processor.
2. Wait for `MSRC001_0071[CurCpuFid, CurCpuDid] = [CpuFid, CpuDid]` from `MSRC001_00[6B:64]` indexed by `D18F3xDC[HwpstateMaxVal]`.
3. Step 2 must be completed on all cores prior to executing step 4 since a compute unit transitions to the highest performance P-state requested on either core.
4. Write 0 to `MSRC001_0062[PstateCmd]` on all cores in the processor.
5. Wait for `MSRC001_0071[CurCpuFid, CurCpuDid] = [CpuFid, CpuDid]` from `MSRC001_00[6B:64]` indexed by `MSRC001_0071[CurPstateLimit]`.
6. If `MSRC001_0071[CurPstateLimit] != D18F3xDC[HwpstateMaxVal]`, wait for `MSRC001_0071[CurCpuVid] = [CpuVid]` from `MSRC001_00[6B:64]` indexed by `MSRC001_0071[CurPstateLimit]`.
7. Wait for `MSRC001_0063[CurPstate] = MSRC001_0061[CurPstateLimit]`.

2.5.3.1.9.2 Core Minimum P-state Transition Sequence After Warm Reset

1. Write 0 to `MSRC001_0062[PstateCmd]` on all cores in the processor.
2. Wait for `MSRC001_0071[CurCpuFid, CurCpuDid] = [CpuFid, CpuDid]` from `MSRC001_00[6B:64]` indexed by `MSRC001_0071[CurPstateLimit]`.
3. Write `MSRC001_0061[PstateMaxVal]` to `MSRC001_0062[PstateCmd]` on all cores in the processor.
4. Wait for `MSRC001_0071[CurCpuFid, CurCpuDid] = [CpuFid, CpuDid]` from `MSRC001_00[6B:64]` indexed by `D18F3xDC[HwpstateMaxVal]`.
5. If `MSRC001_0071[CurPstateLimit] != MSRC001_0071[CurPstate]`, wait for `MSRC001_0071[CurCpuVid] = [CpuVid]` from `MSRC001_00[6B:64]` indexed by `D18F3xDC[HwpstateMaxVal]`.
6. Wait for `MSRC001_0063[CurPstate] = MSRC001_0062[PstateCmd]`.

2.5.3.1.9.3 ACPI Processor P-state Objects

Processor performance control is implemented through the `_PCT`, `_PSS` and `_PSD` objects in ACPI 2.0 and later revisions. The presence of these objects indicates to the OS that the platform and processor are capable of supporting multiple performance states. Processor performance states are not supported with ACPI 1.0b. BIOS must provide the `_PCT`, `_PSS`, and `_PSD` objects, and define other ACPI parameters to support operating systems that provide native support for processor P-state transitions.

The following rules apply to BIOS generated ACPI objects in multi-core systems. Refer to the appropriate ACPI specification (<http://www.acpi.info>) for additional details:

- All cores must expose the same number of performance states to the OS.
- The respective performance states displayed to the OS for each core must have identical performance and power-consumption parameters (e.g. P0 on core 0 must have the same performance and power-consumptions parameters as P0 on core 1, P1 on core 0 must have the same parameters as P1 on core 1, however P0 can be different than P1).
- Performance state objects must be present under each processor object in the system.

2.5.3.1.9.3.1 `_PCT` (Performance Control)

BIOS must declare the performance control object parameters as functional fixed hardware. This definition indicates the processor driver understands the architectural definition of the P-state interface associated with `CPUID Fn8000_0007_EDX[Hwpstate]=1`.

- `Perf_Ctrl_Register` = Functional Fixed Hardware
- `Perf_Status_Register` = Functional Fixed Hardware

2.5.3.1.9.3.2 **_PSS (Performance Supported States)**

A unique `_PSS` entry is created for each non-boosted P-state. The value contained in the `_PSS Control` field is written to `MSRC001_0062 [P-state Control]` to request a P-state change to the CoreFreq of the associated `_PSS` object. The value contained in `MSRC001_0063 [P-state Status]` can be used to identify the `_PSS` object of the current P-state request by equating `MSRC001_0063[CurPstate]` to the value of the Status field. See 2.5.3.1 [Core P-states].

BIOS loops through each of `MSRC001_00[6B:64]` applying the following formulas to create the fields for the `_PSS` object for each valid P-state (see `MSRC001_00[6B:64][PstateEn]`). BIOS skips over any P-state MSRs that specify boost P-states (see `D18F4x15C[NumBoostStates]`).

- CoreFreq (MHz) = Calculated using the formula for `CoreCOF`.
- Power (mW) = `MSRC001_00[6B:64][CpuVid]` voltage * `MSRC001_00[6B:64][IddValue]` current * 1000.
- TransitionLatency (us) and BusMasterLatency (us):
 - If `MSRC001_00[6B:64][CpuFid]` is the same for all enabled P-states (see `MSRC001_00[6B:64][PstateEn]`) and all boosted P-states:
 - TransitionLatency = BusMasterLatency = (15 steps * `D18F3xD4[PowerStepDown]` time * 1000 us/ns) + (15 steps * `D18F3xD4[PowerStepUp]` time * 1000 us/ns)
 - Else if `MSRC001_00[6B:64][CpuFid]` is different for any enabled (see `MSRC001_00[6B:64][PstateEn]`) or boost P-states:
 - TransitionLatency = BusMasterLatency = (15 steps * `D18F3xD4[PowerStepDown]` time * 1000 us/ns) + `D18F3xA0[PIILockTime]` time + (15 steps * `D18F3xD4[PowerStepUp]` time * 1000 us/ns)
- Example:
 - `MSRC001_00[6B:64][CpuFid]` is not the same for all P-states
 - `D18F3xD4[PowerStepDown]` = `D18F3xD4[PowerStepUp]` = 8h (50 ns/step)
 - `D18F3xA0[PIILockTime]` = 001b (2 us)
 - TransitionLatency = BusMasterLatency = (15 steps * 50 ns/step / 1000 us/ns) + 2us + (15 steps * 50 ns/step / 1000 us/ns) = 3.5 us (round up to 4 us)
- Control/Status:
 - The highest performance non-boosted P-state must have the `_PSS control` and status fields programmed to 0.
 - Any lower performance non-boosted P-states must have the `_PSS control` and status fields programmed in ascending order.

2.5.3.1.9.3.3 **_PPC (Performance Present Capabilities)**

The `_PPC` object is optional. Refer to the ACPI specification for details on use and content.

2.5.3.1.9.3.4 **_PSD (P-state Dependency)**

AMD recommends the ACPI 3.0 `_PSD` object be generated for each core as follows to cause the cores to transition between P-states together:

- NumberOfEntries = 5.
- Revision = 0.
- Domain = 0.
- CoordType = FCh. (SW_ALL)
- NumProcessors = `CPUID Fn8000_0008_ECX[NC]` + 1.

A vendor may optionally choose to generate _PSD object to allow cores to transition between P-states independently as follows:

- NumberOfEntries = 5.
- Revision = 0.
- Domain = [CPUID Fn0000_0001_EBX\[LocalApicId\[7:1\]\]](#).
- CoordType = FEh. (HW_ALL)
- NumProcessors = 2.

BIOS provides an option to choose between either _PSD definition.

2.5.3.1.9.4 Fixed ACPI Description Table (FADT) Entries

Declare the following FADT entries:

- PSTATE_CNT = 00h.
- DUTY_WIDTH = 00h.

2.5.3.1.9.5 XPSS (Microsoft® Extended PSS) Object

Some Microsoft® operating systems require an XPSS object to make P-state changes function properly. A BIOS that implements an XPSS object has special requirements for the _PCT object. See the Microsoft *Extended PSS ACPI Method Specification* for the detailed requirements to implement these objects.

2.5.3.2 Core C-states

C-states are processor power states. C0 is the operational state in which instructions are executed. All other C-states are low-power states in which instructions are not executed. When coming out of warm and cold reset, the processor is transitioned to the C0 state.

2.5.3.2.1 C-state Names and Numbers

C-states are often referred to by an alphanumeric naming convention, C1, C2, C3, etc. The mapping between ACPI defined C-states and AMD specified C-states is not direct. AMD specified C-states are referred to as IO-based C-states. Up to three IO-based C-states are supported, IO-based C-state 0, 1, and 2. The IO-based C-state index corresponds to the offset added to [MSRC001_0073\[CstateAddr\]](#) to initiate a C-state request. See [2.5.3.2.2 \[C-state Request Interface\]](#). The actions taken by the processor when entering a low-power C-state are configured by software. See [2.5.3.2.3 \[C-state Actions\]](#) for information about AMD specific actions.

2.5.3.2.2 C-state Request Interface

C-states are dynamically requested by software and are exposed through ACPI objects (see [2.5.3.2.6 \[ACPI Processor C-state Objects\]](#)). C-states can be requested on a per-core basis. Software requests a C-state change in one of two ways:

- Reading from an IO address: The IO address must be the address specified by [MSRC001_0073\[CstateAddr\]](#) plus an offset of 0 through 7. The processor always returns 0 for this IO read. Offsets 2 through 7 result in an offset of 2.
- Executing the HLT instruction. This is equivalent to reading from the IO address specified by [D18F4x128\[HaltCstateIndex\]](#).

When software requests a C-state transition, hardware evaluates any frequency and voltage domain dependencies and determines which C-state actions to execute. See [2.5.2 \[Frequency and Voltage Domain Dependen-](#)

cies] and 2.5.3.2.3 [C-state Actions].

2.5.3.2.3 C-state Actions

A core takes one of several different possible actions based upon a C-state change request from software. The C-state action fields are defined in [D18F4x11\[C:8\]](#) and [D18F4x11C](#).

2.5.3.2.3.1 C-state Probes and Cache Flushing

If probes occur after a core enters a non-C0 state, and the caches are not flushed by hardware, the core clock may be ramped back up to the C0 frequency to service the probes, as specified by [D18F3x\[84:80\]\[CpuPrbEn\]](#) or [D18F4x11\[C:8\]/D18F4x11C\[*CpuPrbEn\]](#).

If a core enters a non-C0 state and cache flush is enabled (see [D18F3xDC\[CacheFlushOnHaltCtl\]](#) and [D18F4x11\[C:8\]\[CacheFlushEn\]](#)), a timer counts down for a programmable period of time as specified by [D18F3xDC\[CacheFlushOnHaltTmr\]](#) or [D18F4x11\[C:8\]\[CacheFlushTmrSel\]](#). When the timer expires, the core flushes its L1 and L2 caches to DRAM and the core clocks are ramped down to a divisor specified by [D18F3xDC\[CacheFlushOnHaltCtl\]](#). The timer is reset if the core exits the C-state for any reason. See 2.5.3.2.4.3 [Cache Flush On Halt Saturation Counter].

Once a core flushes its caches, probes are no longer sent to that core. This improves probing performance for cores that are in C0.

2.5.3.2.3.2 Core C1 (CC1) State

When a core enters the CC1 state, its clock ramps down to the frequency specified by [D18F4x11\[C:8\]\[ClkDivisorCstAct\]](#).

2.5.3.2.3.3 Core C6 (CC6) State

A core can gate off power to its internal logic when it enters any non-C0 state. This power gated state is known as CC6. In order to enter CC6, hardware first enters CC1 and then flushes the caches (see 2.5.3.2.3.1 [C-state Probes and Cache Flushing]) before checking [D18F4x11\[C:8\]\[PwrGateEnCstAct\]](#). Power gating reduces the amount of power consumed by the core. VDD voltage is not reduced when a core is in CC6. The following sequence occurs when a core enters the CC6 state:

1. If [MSRC001_0071\[CurPstate\]](#) < [D18F3xA8\[PopDownPstate\]](#), transition the core P-state to [D18F3xA8\[PopDownPstate\]](#).
2. L1 and L2 caches are flushed to DRAM. See 2.5.3.2.3.1 [C-state Probes and Cache Flushing].
3. Internal core state is saved to DRAM.
4. Power is removed from the core and the core PLL/voltage regulator is powered down as specified by [D18F5x128\[CC6PwrDwnRegEn\]](#).

All of the following must be true in order for a core to be placed into CC6:

- [D18F4x11\[C:8\]\[CacheFlushEn\]](#)=1 for the corresponding C-state action field.
- [D18F4x11\[C:8\]\[CacheFlushTmrSel\]](#) != 11b for the corresponding C-state action field.
- [D18F4x11\[C:8\]\[PwrGateEnCstAct\]](#)=1 for the corresponding C-state action field.
- [D18F2x118\[CC6SaveEn\]](#)=1.
- [D18F2x118\[LockDramCfg\]](#)=1 .
- The CC6 storage area in DRAM is configured. See 2.9.8 [DRAM CC6/PC6 Storage].

The events which cause a core to exit the CC6 state are specified in [2.5.3.2.5 \[Exiting C-states\]](#).

If a warm reset occurs while a core is in CC6, all MCA registers in the core shown in [Table 56](#) are cleared to 0. See [2.15.1 \[Machine Check Architecture\]](#).

The time required to enter and exit CC6 is directly proportional to the core P-state frequency. Slower core frequencies require longer entry and exit times. Latency issues may occur with core P-state frequencies less than 800MHz.

2.5.3.2.3.4 Package C6 (PC6) State

When all cores enter a non-C0 state, VDD can be reduced to a non-operational voltage that does not retain core state. This state is known as PC6 and reduces the amount of static and dynamic power consumed by all cores. The following actions are taken by hardware prior to PC6 entry:

1. If [MSRC001_0071\[CurPstate\]](#) < [D18F3xA8\[PopDownPstate\]](#), transition the core P-state to [D18F3xA8\[PopDownPstate\]](#).
2. For all cores not in CC6, L1 and L2 caches are flushed to DRAM. See [2.5.3.2.3.1 \[C-state Probes and Cache Flushing\]](#).
3. For all cores not in CC6, internal core state is saved to DRAM.
4. VDD is transitioned to the VID specified by [D18F5x128\[PC6Vid\]](#).
5. If the core PLLs are not powered down during CC6 entry (see [2.5.3.2.3.3 \[Core C6 \(CC6\) State\]](#)), then they are powered down as specified by [D18F5x128\[PC6PwrDwnRegEn\]](#).

All of the following must be true on all cores in order for a package to be placed into PC6:

- [D18F4x11\[C:8\]\[CacheFlushEn\]=1](#) for the corresponding C-state action field
- [D18F4x11\[C:8\]\[CacheFlushTmrSel\] != 11b](#) for the corresponding C-state action field.
- [D18F4x11\[C:8\]\[PwrOffEnCstAct\]=1](#) for the corresponding C-state action field.
- [D18F2x118\[CC6SaveEn\]=1](#).
- [D18F2x118\[LockDramCfg\]=1](#).
- [MSRC001_0015\[HltXSpCycEn\]=1](#).
- Timer Tick Monitor does not disallow PC6 entry. See [2.5.3.2.4.2 \[Timer Tick Monitor\]](#).

2.5.3.2.4 C-state Request Monitors

Deeper C-states have higher entry and exit latencies but provide greater power savings than shallower C-states. To help balance the performance and power needs of the system, the processor can limit access to specific C-states in certain scenarios.

2.5.3.2.4.1 FCH Messaging

The FCH can be notified when the processor transitions package C-states. See the following:

- [D18F4x128\[CstateMsgDis\]](#).
- [D18F5x178\[CstateFusionDis\]](#).
- [MSRC001_0015\[HltXSpCycEn\]](#).

If the processor sends a message when entering PC6, the FCH sends a return message notifying the processor whether PC6 is allowed. The processor can ignore the FCH response as specified by [D18F5x178\[CstateFusionHsDis\]](#). A C-state message indicating the package C-state entered is sent to the FCH as specified by [D18F5x178\[CstateThreeWayHsEn\]](#).

2.5.3.2.4.2 Timer Tick Monitor

The timer tick monitor (enabled using [D18F4x124\[IntMonPC6En\]](#)) tracks interrupts that are delivered from the FCH on a regular interval, generally referred to as timer tick interrupts. When software changes the period of any timers in the FCH, the FCH reports the new period to the processor. The timer tick monitor cannot be used to track periodic local APIC timer interrupts.

The monitor compares the timer tick period to the threshold defined by [D18F4x124\[IntMonPC6Limit\]](#). If the period is less than or equal to the threshold, access to PC6 is disallowed.

2.5.3.2.4.3 Cache Flush On Halt Saturation Counter

A cache flush success monitor tracks the success rate of cache flush timer expirations relative to the core exiting a C-state. Based on the success rate, caches may be flushed immediately without waiting for the cache flush timer to expire. See [D18F4x128\[CacheFlushSucMonThreshold\]](#). When the core resumes normal execution, the caches refill as normal.

2.5.3.2.5 Exiting C-states

The following events may cause the processor to exit a non-C0 C-state and return to C0:

- INTR
- NMI
- SMI
- INIT
- RESET_L assertion

If an INTR is received while a core is in a non-C0 C-state, the state of [EFLAGS\[IF\]](#) and the mechanism used to enter the non-C0 C-state determine the actions taken by the processor.

- Entry via HLT, [EFLAGS\[IF\]](#)==0: The interrupt does not wake up the core.
- Entry via HLT, [EFLAGS\[IF\]](#)==1: The interrupt wakes the core and the core begins execution at the interrupt service routine.
- Entry via IO read, [EFLAGS\[IF\]](#)==0: The interrupt wakes the core and the core begins execution at the instruction after the IN instruction that was used to enter the non-C0 C-state.
- Entry via IO read, [EFLAGS\[IF\]](#)==1: The interrupt wakes the core and the core begins execution at the interrupt service routine.

2.5.3.2.6 ACPI Processor C-state Objects

Processor power control is implemented through the [_CST](#) object in ACPI 2.0 and later revisions. The presence of the [_CST](#) object indicates to the OS that the platform and processor are capable of supporting multiple power states. BIOS must provide the [_CST](#) object and define other ACPI parameters to support operating systems that provide native support for processor C-state transitions. See [2.5.3.2.6.1 \[_CST\]](#) and [2.5.3.2.6.2 \[_CSD\]](#).

The [_CST](#) object is not supported with ACPI 1.0b. BIOS should provide FADT entries to support operating systems that are not compatible with ACPI 2.0 and later revisions. See [2.5.3.2.6.4 \[Fixed ACPI Description Table \(FADT\) Entries\]](#).

2.5.3.2.6.1 [_CST](#)

The [_CST](#) object should be generated for each core as follows:

- Count = 1.

- Register = [MSRC001_0073](#)[CstateAddr] + 1.
- Type = 2.
- Latency = 100.
- Power = 0.

2.5.3.2.6.2 [_CSD](#)

The [_CSD](#) object should be generated for each core as follows:

- NumberOfEntries = 6.
- Revision = 0.
- Domain = [CPUID Fn0000_0001_EBX](#)[LocalApicId[7:1]].
- CoordType = FEh. (HW_ALL)
- NumProcessors = 2.
- Index = 0.

2.5.3.2.6.3 [_CRS](#)

BIOS must declare in the root host bridge [_CRS](#) object that the IO address range from [MSRC001_0073](#)[CstateAddr] to [MSRC001_0073](#)[CstateAddr]+7 is consumed by the host bridge.

2.5.3.2.6.4 Fixed ACPI Description Table (FADT) Entries

Declare the following FADT entries:

- P_LVL2_LAT = 100.
- P_LVL3_LAT = 1001.
- FLAGS.PROC_C1 = 1.
- FLAGS.P_LVL2_UP = 1.

Declare the following P_BLK entries:

- P_LVL2 = [MSRC001_0073](#)[CstateAddr] + 1.
- P_LVL3 = 0.

BIOS must declare the PSTATE_CNT entry as 00h.

2.5.3.2.7 BIOS Requirements for Initialization

1. Initialize [MSRC001_0073](#)[CstateAddr] with an available IO address. See 2.5.3.2.6.3 [[_CRS](#)].
2. Initialize [D18F4x11](#)[C:8].
3. Generate ACPI objects as described in 2.5.3.2.6 [[ACPI Processor C-state Objects](#)].

2.5.3.3 Effective Frequency

The effective frequency interface allows software to discern the average, or effective, frequency of a given core over a configurable window of time. This provides software a measure of actual performance rather than forcing software to assume the current frequency of the core is the frequency of the last P-state requested. This can be useful when the P-state is limited by HTC, [D18F3x68](#), SBI, or CPB.

The following procedure calculates effective frequency using [MSR0000_00E7](#) [Max Performance Frequency Clock Count (MPERF)] and [MSR0000_00E8](#) [Actual Performance Frequency Clock Count (APERF)]:

1. At some point in time, write 0 to both MSRs.
2. At some later point in time, read both MSRs.

- Effective frequency = (value read from MSR0000_00E8 / value read from MSR0000_00E7) * P0 frequency using software P-state numbering.

Additional notes:

- The amount of time that elapses between steps 1 and 2 is determined by software.
- It is software's responsibility to disable interrupts or any other events that may occur in between the write of MSR0000_00E7 and the write of MSR0000_00E8 in step 1 or between the read of MSR0000_00E7 and the read of MSR0000_00E8 in step 2.
- The behavior of MSR0000_00E7 and MSR0000_00E8 may be modified by MSRC001_0015[EffFreqCntMwait].
- The effective frequency interface provides +/- 50MHz accuracy if the following constraints are met:
 - Effective frequency is read at most one time per millisecond.
 - When reading or writing MSR0000_00E7 and MSR0000_00E8 software executes only MOV instructions, and no more than 3 MOV instructions, between the two RDMSR or WRMSR instructions.
- MSR0000_00E7 and MSR0000_00E8 are invalid if an overflow occurs.

2.5.4 NB Power Management

2.5.4.1 NB P-states

The processor supports up to four NB P-states (NBP0 through NBP3), specified in D18F5x1[6C:60]. Each NB P-state consists of the following:

- Enable: D18F5x1[6C:60][NbPstateEn].
- NCLK frequency: D18F5x1[6C:60][NbFid, NbDid].
- VDDNB voltage: D18F5x1[6C:60][NbVid].
- Memory P-state: D18F5x1[6C:60][MemPstate]. See 2.5.7.1 [Memory P-states].

Although four NB P-states are defined, only two NB P-states are used at any given time, specified by D18F5x170[NbPstateHi, NbPstateLo]. Hardware dynamically changes which two of the four NB P-states are in use based on GPU activity as follows:

- When the GPU is active:
 - The high NB P-state is specified by D0F0xBC_x1F5F8[DpmXNbPsHi].
 - The low NB P-state is specified by D0F0xBC_x1F5F8[DpmXNbPsLo].
- When the GPU is idle, and the timer specified by D0F0xBC_x1F5F8[Hysteresis] has expired:
 - The high NB P-state is specified by D0F0xBC_x1F5F8[Dpm0PgNbPsHi].
 - The low NB P-state is specified by D0F0xBC_x1F5F8[Dpm0PgNbPsLo].

In addition, hardware forces the NB P-state to the active high or low NB P-state based on the GPU activity level. These decisions are all made at the rate specified by D0F0xBC_x1F638[NbDpmPeriod].

Out of cold reset, the NB P-state is specified by D18F5x174[StartupNbPstate] and D18F3xA0[CofVidProg]. The current NB P-state is specified by D18F5x174[CurNbFid, CurNbDid, CurNbVid].

2.5.4.1.1 NB P-state Transitions

Hardware selects whether to use the high or low NB P-state based on several criteria as follows:

- Core P-state:

- MSRC001_00[6B:64][NbPstate].
- D18F5x170[NbPstateThreshold].
- GPU activity:
 - The GPU driver selects levels of GPU activity that force the NB P-state to either the high or low state or allow either NB P-state.
- Hysteresis timer:
 - D18F5x170[NbPstateHiRes, NbPstateLoRes].
- The following configuration registers:
 - D18F5x170[SwNbPstateLoDis, NbPstateDisOnP0].
 - MSRC001_0071[NbPstateDis].

Note: BIOS should set D18F5x170[SwNbPstateLoDis] to 1b, to be cleared by the GPU driver when it loads.

Once hardware determines that an NB P-state transition is necessary, hardware executes the following sequence:

1. If transitioning from the low NB P-state to the high NB P-state, transition VDDNB voltage.
2. If the GPU is enabled as specified by D18F5x178[SwGfxDis], wait for display buffer to fill.
3. Quiesce all active cores.
4. If the internal GPU is enabled as specified by D18F5x178[SwGfxDis] and D18F5x178[Dis2ndGnbAllowPsWait]==0, wait for display buffer to fill.
5. Stop memory traffic and place DRAM into self-refresh.
6. Transition NCLK frequency.
7. Update NB P-state specific DRAM settings within hardware, see D18F2x210_dct[1:0]_nbp[3:0].
8. Take DRAM out of self-refresh and allow memory traffic.
9. Wake up cores.
10. If transitioning from the high NB P-state to the low NB P-state, transition VDDNB voltage.

2.5.4.1.2 BIOS NB P-state Configuration

2.5.4.1.2.1 NB P-state COF and VID Synchronization After Warm Reset

BIOS performs the following sequence on one core. All APs must be initialized prior to running the steps in this section. See 2.3.2 [AP initialization]. This is done after any warm reset and before 2.9.5 [DCT/DRAM Initialization and Resume]:

1. Temp1=D18F5x170[SwNbPstateLoDis].
2. Temp2=D18F5x170[NbPstateDisOnP0].
3. Temp3=D18F5x170[NbPstateThreshold].
4. Temp4=D18F5x170[NbPstateGnbSlowDis].
5. If MSRC001_0070[NbPstate]=0, go to step 6. If MSRC001_0070[NbPstate]=1, go to step 11.
6. Write 1 to D18F5x170[NbPstateGnbSlowDis].
7. Write 0 to D18F5x170[SwNbPstateLoDis, NbPstateDisOnP0, NbPstateThreshold].
8. Wait for D18F5x174[CurNbPstate] = D18F5x170[NbPstateLo] and D18F5x174[CurNbFid, CurNbDid]=[NbFid, NbDid] from D18F5x1[6C:60] indexed by D18F5x170[NbPstateLo].
9. Set D18F5x170[SwNbPstateLoDis]=1.
10. Wait for D18F5x174[CurNbPstate] = D18F5x170[NbPstateHi] and D18F5x174[CurNbFid, CurNbDid]=[NbFid, NbDid] from D18F5x1[6C:60] indexed by D18F5x170[NbPstateHi]. Go to step 15.
11. Write 1 to D18F5x170[SwNbPstateLoDis].
12. Wait for D18F5x174[CurNbPstate] = D18F5x170[NbPstateHi] and D18F5x174[CurNbFid, CurNbDid]=[NbFid, NbDid] from D18F5x1[6C:60] indexed by D18F5x170[NbPstateHi].
13. Write 0 to D18F5x170[SwNbPstateLoDis, NbPstateDisOnP0, NbPstateThreshold].

14. Wait for `D18F5x174[CurNbPstate] = D18F5x170[NbPstateLo]` and `D18F5x174[CurNbFid, CurNbDid]=[NbFid, NbDid]` from `D18F5x1[6C:60]` indexed by `D18F5x170[NbPstateLo]`.
15. Set `D18F5x170[SwNbPstateLoDis]=Temp1`, `D18F5x170[NbPstateDisOnP0]=Temp2`, `D18F5x170[NbPstateThreshold]=Temp3`, and `D18F5x170[NbPstateGnbSlowDis]=Temp4`.

2.5.4.1.2.2 NB P-state Voltage Adjustment for MEMCLK

During 2.9.5 [DCT/DRAM Initialization and Resume], BIOS updates the NB P-state voltages based on the target MEMCLK for M0 and M1 using the following steps. See 2.9.5.2 [NB P-state Specific Configuration].

1. Determine the target MEMCLK for M0 and M1 (see 2.9.5 [DCT/DRAM Initialization and Resume] and 2.5.7.1 [Memory P-states]).
2. `MemClkVidHi` = read `D0F0xBC_xE0104168` through `D0F0xBC_xE0104170` to find the VID code corresponding to the M0 MEMCLK. If the M0 MEMCLK is not found, use the next higher defined MEMCLK as the target. If the M0 MEMCLK is greater than frequency associated with `D0F0xBC_xE0104170[MemClkVid8]`, use `D0F0xBC_xE0104170[MemClkVid8]` as the target.
3. `MemClkVidLo` = read `D0F0xBC_xE0104168` through `D0F0xBC_xE0104170` to find the VID code corresponding to the M1 MEMCLK. If the M1 MEMCLK is not found, use the next higher defined MEMCLK as the target. If the M1 MEMCLK is greater than frequency associated with `D0F0xBC_xE0104170[MemClkVid8]`, use `D0F0xBC_xE0104170[MemClkVid8]` as the target.
4. For each NB P-state from `NBP0` through `D18F5x170[NbPstateMaxVal]`:
 - If `((D18F5x1[6C:60][MemPstate] == 0) && (MemClkVidHi voltage > D18F5x1[6C:60][NbVid] voltage))`:
 - Program `D18F5x1[6C:60][NbVid] == MemClkVidHi`.
 - If `((D18F5x1[6C:60][MemPstate] == 1) && (MemClkVidLo voltage > D18F5x1[6C:60][NbVid] voltage))`:
 - Program `D18F5x1[6C:60][NbVid] == MemClkVidLo`.
5. Program = `MemClkVidHi`.

2.5.4.1.2.3 NB P-state Configuration for Runtime

Please see your AMD representative for details.

2.5.4.2 NB C-states

NB C-states are package-level actions that occur only when all cores enter a non-C0 state (see 2.5.3.2 [Core C-states]). The NB C-state actions are:

- DRAM self-refresh (see 2.5.7.2 [DRAM Self-Refresh]):
 - Enable bit: `D18F4x11[C:8][SelfRefr]`.
 - Entry requirements:
 - No outstanding GPU traffic or traffic from a link.
 - Exit conditions (any of the following must be true):
 - The local APIC timer expires. See 2.4.8.1 [Local APIC].
 - New GPU traffic or traffic from a link.
 - A P-state limit update (see 2.5.3.1.5 [Core P-state Limits]) causes the most restrictive P-state limit to become a higher number than the current P-state for any core in CC1.
- NB power gating:
 - Enable bit: `D18F4x11[C:8][NbPwrGate]`.
 - Entry requirements (all of the following must be true):

- No outstanding GPU traffic or traffic from a link.
- All cores are in CC6.
- DRAM is either in or entering self-refresh.
- Exit conditions (any of the following must be true):
 - The local APIC timer expires. See 2.4.8.1 [Local APIC].
 - New GPU traffic or traffic from a link.
 - A P-state limit update (see 2.5.3.1.5 [Core P-state Limits]) causes the most restrictive P-state limit to become a higher number than the current P-state for any core in CC1.
- NB clock gating:
 - Enable bit: `D18F4x11[C:8][NbClkGate]`.
 - Entry requirements:
 - No outstanding GPU traffic or traffic from a link.
 - Exit conditions (any of the following must be true):
 - The local APIC timer expires. See 2.4.8.1 [Local APIC].
 - New GPU traffic or traffic from a link.
 - A P-state limit update (see 2.5.3.1.5 [Core P-state Limits]) causes the most restrictive P-state limit to become a higher number than the current P-state for any core in CC1.

When entering NB C-states, the actions are taken in the following order:

1. DRAM self-refresh.
2. NB clock gating.
3. NB power gating.

When exiting NB C-states, the actions are taken in the following order:

1. NB power gating.
2. NB clock gating.
3. DRAM self-refresh.

2.5.5 P-state Bandwidth Requirements

- The frequency relationship of (core COF / NB COF) ≤ 6 must be maintained for all supported P-state combinations. E.g., a core COF of 4.0 GHz could not be combined with a NB COF of 0.6 GHz; the NB COF would have to be 0.7 GHz or greater.
- All core P-states are required to be defined such that (NB COF/core COF) ≤ 32 , for all NB/core P-state combinations. E.g., if the NB COF is 4.8 GHz then the core COF must be no less than 150 MHz.
- All core P-states must be defined such that `CoreCOF` ≥ 900 MHz.
- All core P-states must be defined such that `MSRC001_00[6B:64][CpuFid]` $\leq 22h$.
- All NB P-states must be defined such that `D18F5x1[6C:60][NbFid]` $\leq 2Eh$.
- `NBCOF` $> 1.25 * \text{MEMCLK}$ frequency.
- `NBCOF` $\geq 700\text{Mhz}$.
- If `NBCOF` $= 700\text{Mhz}$ then MEMCLK must be 333Mhz.

2.5.6 GPU and Root Complex Power Management

2.5.6.1 Dynamic Power Management (DPM)

The processor supports dynamic SCLK and LCLK frequency changes along with VDDNB voltage change requests, known as Dynamic Power Management (DPM). Once initialized, hardware dynamically monitors processor utilization and adjusts the frequencies and voltage based on that utilization. For both LCLK and SCLK DPM, higher numbered states represent higher performance and lower numbered states represent lower

performance. All family 15h model 10h-1Fh processors support DPM.

2.5.6.1.1 Activity Monitors

The processor contains two activity monitors, one for 2.5.6.1.2 [SCLK DPM] and one for 2.5.6.1.3 [LCLK DPM]. Each activity monitor tracks the usage level of different processor subcomponents. A binary signal from each subcomponent is used to determine whether that subcomponent is busy. On each SCLK or LCLK cycle, the respective activity monitors sample the signal from each unmasked subcomponent (see D0F0xBC_xE0000120).

The output of the activity monitor is then used to determine whether the DPM state should be changed. See 2.5.6.1.2 [SCLK DPM] and 2.5.6.1.3 [LCLK DPM].

2.5.6.1.2 SCLK DPM

SCLK DPM consists of up to 8 states. Any number of states up through 8 may be used and there is no requirement that the states be contiguous.

When SCLK DPM is enabled, the DPM state is transitioned to D0F0xBC_x1F100[ScIkDpmBootState] and the activity monitor begins calculating activity levels (see 2.5.6.1.1 [Activity Monitors]). Each time a new activity value is calculated, it is placed into a loop filter and delta-sigma modulator. At that point, the SCLK DPM monitor determines whether to transition to next higher numbered (higher performance) valid state or the next lower numbered (lower performance) valid state.

SCLK DPM is enabled by setting D0F0xBC_x1F100[ScIkDpmEn].

During runtime, SCLK DPM parameters are programmed by the graphics driver. To support the driver, BIOS creates a data structure in memory containing information about the processor. Please see your AMD representative for more information.

2.5.6.1.3 LCLK DPM

LCLK DPM is functionally identical to 2.5.6.1.2 [SCLK DPM] except for the locations of the following parameters:

- Valid bit: D0F0xBC_x1F2[E0:00:step20][StateValid].
- Voltage change hysteresis threshold: D0F0xBC_x1F2[E0:00:step20][LowVoltageReqThreshold].
- Divisor: D0F0xBC_x1F2[E0:00:step20][ScIkDivider].
- VID: D0F0xBC_x1F2[E0:00:step20][VID].
 - See 2.5.2.2 [Dependencies Between Subcomponents on VDDNB].
- State change hysteresis thresholds: D0F0xBC_x1F2[E8:08:step20][HysteresisUp, HysteresisDown].
- Activity thresholds: D0F0xBC_x1F2[F0:10:step20][ActivityThreshold].
- Residency counter: See D0F0xBC_x1F2[E8:08:step20][ResidencyCounter].

LCLK DPM is enabled by setting D0F0xBC_x1F300[LclkDpmEn]. LCLK DPM voltage changes are enabled using D0F0xBC_x1F300[VoltageChgEn]. When LCLK DPM is first enabled, the DPM state is transitioned to D0F0xBC_x1F300[LclkDpmBootState].

2.5.6.2 GPU and Root Complex Power Gating

Several subcomponents of the GPU and root complex can be power gated when not in use.

- **GPU:** GPU power gating is initialized and enabled by software (see [GpuEnabled](#) and [D0F0x7C\[ForceIntGfxDisable\]](#)). Once initialized and enabled, the GPU is power gated by hardware when inactive and is ungated by hardware when needed. When internal GPU is disabled by BIOS, BIOS is responsible for power gating the GPU.
- **GMC:** GMC power gating is initialized and enabled by software. Once initialized and enabled, GMC is power gated by hardware when inactive and is ungated by hardware when needed. GMC's state is saved internally. If the internal GPU is disabled, either by hardware (fusing) or by software, software is responsible for power gating GMC.
- **UVD:** UVD is power gated by software when not in use and is ungated by software when needed. UVD's internal state is not saved and UVD goes through an internal reset when power is restored.
- **VCE:** VCE is power gated by software when not in use and is ungated by software when needed. VCE's internal state is not saved and VCE goes through an internal reset when power is restored.
- **DCE:** DCE is power gated by software when there is no display connected. DCE's internal state is not saved and DCE goes through an internal reset when power is restored.
- **PCIe:** The Gfx link core can be power gated when it is not in use. In addition, the individual phys on the TX and RX sides of each link can be power gated when the links are not connected. During POST and runtime, several software components inform the SMU whether the Gfx link core or the link phys are in use. Hardware dynamically power gates or ungates the Gfx link core and the link phys as needed.

2.5.7 DRAM Power Management

2.5.7.1 Memory P-states

The processor supports up to 2 memory P-states, M0 and M1. Each memory P-state consists of the following:

- MEMCLK frequency: [D18F2x94_dct\[1:0\]\[MemClkFreq\]](#) and [D18F2x2E0_dct\[1:0\]\[M1MemClkFreq\]](#).
- A set of frequency dependent timing and configuration registers: See [2.9.1 \[DCT Configuration Registers\]](#).

All valid memory P-states are associated with a specific NB P-state, as specified by [D18F5x1\[6C:60\]\[MemPstate\]](#). When hardware transitions to a new NB P-state, the memory P-state is transitioned to that specified by the new NB P-state.

Out of cold reset the current memory P-state is M0. The P-state value specified by [D18F5x1\[6C:60\]\[MemPstate\]](#) of the NB P-state indexed by [D18F5x174\[StartupNbPstate\]](#) is invalid. Support for dynamic memory P-state changes is indicated by [D18F3xE8\[MemPstateCap\]=1](#) and one or more [D18F5x1\[6C:60\]\[MemPstate\]=1](#); otherwise M0 is used by hardware for configuration purposes.

2.5.7.2 DRAM Self-Refresh

DRAM is placed into self-refresh on S3 entry (see [2.5.8.1.1 \[ACPI Suspend to RAM State \(S3\)\]](#)).

In addition to S3, DRAM is placed into self-refresh in S0 in the following two scenarios:

- NB P-state transitions (see [2.5.4.1 \[NB P-states\]](#)).
- NB C-states (see [2.5.4.2 \[NB C-states\]](#)).

The following requirements must be met before hardware places DRAM into self-refresh:

- No pending traffic.
- One of the following is true:
 - The GPU is idle and the internal display buffer is full.
 - The internal GPU is disabled.

Once the above requirements are met, hardware places DRAM into self-refresh.

Early self-refresh occurs when DRAMs are placed in self-refresh before expiration of the cache flush timer. See [D18F4x11\[C:8\]\[SelfRefrEarly\]](#) and [D18F5x128\[SelfRefrEarlyDis\]](#). If early self-refresh is enabled, the DRAMs are taken out of self-refresh to perform the flush operation when the cache flush timer expires and then placed back into self-refresh.

The following are events that cause DRAM to transition out of self-refresh:

- Core transitioning to C0.
- Incoming request from any link or the GMC
- P-state limit update, only in the case when all cores are not in the power gated (CC6) state.

To save additional power, hardware always tristates MEMCLK when entering self-refresh.

2.5.7.3 Stutter Mode

DRAM is most commonly placed in self-refresh due to stutter mode when the internal GPU is in use. The display buffer in the GPU is a combination of a large buffer known as the DMIF (Display Memory Interface FIFO) and a smaller line buffer. The DMIF takes data originating from DRAM, and sends it to the line buffer to draw to the screen. When the data level in the DMIF is full, DRAM is placed in self-refresh, and incoming DRAM requests are queued. As the DMIF drains, it eventually falls below a predefined watermark level, at which point hardware pulls DRAM out of self-refresh and services all the requests in the queue. Once all the requests are complete and the DMIF is full again, a transition back into self-refresh occurs if the stutter mode conditions are still met.

2.5.7.3.1 System BIOS Requirements for Stutter Mode Operation During POST

BIOS creates a data structure in memory containing information about the processor for use by the driver. Please see your AMD representative for more information.

2.5.7.4 EVENT_L

EVENT_L is a level sensitive input to the processor. When asserted, the actions specified by [D18F2xA4](#) are taken. EVENT_L is generally asserted to indicate that a DRAM high temperature condition exists. The minimum assertion time for EVENT_L is 15 ns. The minimum deassertion time for EVENT_L is 15 ns.

- EVENT_L is pulled to VDDIO on the motherboard.
- EVENT_L is ignored while:
 - PWROK is de-asserted.
 - RESET_L is asserted.
- BIOS must ensure that throttling is disabled (see [D18F2xA4\[CmdThrottleMode\]](#)) until DRAM training is complete.

See [2.9.9 \[DRAM On DIMM Thermal Management and Power Capping\]](#).

2.5.8 System Power Management

2.5.8.1 S-states

S-states are ACPI defined sleep states. S0 is the operational state. All other S-states are low-power states in which various voltage rails in the system may or may not be powered. See the ACPI specification for descriptions of each S-state.

2.5.8.1.1 ACPI Suspend to RAM State (S3)

The processor supports the ACPI-defined S3 state. Software is responsible for restoring the state of the processor's registers when resuming from S3. All registers in the processor that BIOS initialized during the initial boot must be restored. The method used to restore the registers is system specific.

During S3 entry, software is responsible for transitioning the processor to Memory Pstate0. See [2.5.7.1 \[Memory P-states\]](#).

During S3 entry, system memory enters self-refresh mode (see [2.5.7.2 \[DRAM Self-Refresh\]](#)). Software is responsible for bringing memory out of self-refresh mode when resuming from S3. To bring memory out of self-refresh mode, see [2.9.5 \[DCT/DRAM Initialization and Resume\]](#).

Many of the systemboard power planes for the processor are powered down during S3. Refer to the *Electrical Data Sheet for AMD Family 15h Models 10h-1Fh Processors* for the following:

- Power plane electrical requirements during S3.
- Power plane sequencing requirements on S3 entry and exit.
- System signal states for both inputs (e.g. PWROK and RESET_L) and outputs (e.g. VID[*], PSI_L bit, THERMTRIP_L, etc.) during S3.
- System signal sequencing requirements on S3 entry and exit.
- System management message sequencing on S3 entry and exit.

2.5.9 Bidirectional Application Power Management (BAPM)

Bidirectional Application Power Management (BAPM) is the algorithm used to contain temperature within a defined limit. BAPM controls several thermal entities. A thermal entity (TE) is any subcomponent of the processor that reports power consumption to BAPM and controls itself to a power limit specified by BAPM. On Trinity, each compute unit and the GPU core are the thermal entities controlled.

On a regular time interval, hardware does the following:

1. Calculates a digital estimate of power consumption for each TE.
2. Converts the power estimates into temperature estimates for each TE.
3. Uses the temperature estimates to assign new power limits to each TE.

Once BAPM has assigned power limits, each TE controls its own frequency and voltage to fit within that limit.

IF (Revision == RL-A1) THEN

2.5.9.1 Hybrid Boost

Hybrid boost is a performance enhancement feature that improves BAPM performance when favorable thermal conditions (like cooler than max-spec ambient, over-the-spec heat sink, etc.) enable the die sensor temperature to be well below the max design temperature limits. When hardware detects such favorable conditions, it

boosts the thermal entities further, beyond what is defined by the baseline BAPM algorithm.

Hybrid Boost is enabled if `D0F0xBC_x1F8EC[HybridBoostNotSupported]` is 0 and `D0F0xBC_x1F428[HybridBoostEn]` is 1.
ENDIF

2.6 Performance Monitoring

The processor includes support for two methods of monitoring processor performance: performance monitor counters and instruction based sampling (IBS). There are two types of performance counters: 2.6.1 [Core Performance Monitor Counters], consisting of one set located in each core of each compute unit and 2.6.2 [NB Performance Monitor Counters].

2.6.1 Core Performance Monitor Counters

The core performance monitor counters are used by software to count specific events that occur in a core of the compute unit. Unless otherwise specified, core performance events count only the activity of the core, not activity caused by the other core of the compute unit. Each core of each compute unit provides six 48-bit performance counters.

MSRC001_020[A,8,6,4,2,0] [Performance Event Select (PERF_CTL[5:0])] specify the events to be monitored and how they are monitored. MSRC001_020[B,9,7,5,3,1] [Performance Event Counter (PERF_CTR[5:0])] are the counters. MSRC001_00[03:00] is the legacy alias for MSRC001_020[6,4,2,0]. MSRC001_00[07:04] is the legacy alias for MSRC001_020[7,5,3,1].

All of the events are specified in 3.22 [Core Performance Counter Events].

Some performance monitor events have a maximum count per clock that exceeds one event per clock. These performance events are called multi-events. Some counters support a greater multi-event count per clock than others. Events that are multi-events will specify the maximum multi-event count per clock. E.g. The number of events logged per cycle can vary from 0 to X. An event that doesn't specify multi-event is implied to be a maximum of 1 event per clock. Undefined results will be produced if an multi-event is selected that exceeds that counters capabilities. The following list specifies the maximum number of multi-events supported by each counter:

- PERF_CTL[0]: 31 multi-event per clock maximum.
- PERF_CTL[1]: 7 multi-event per clock maximum.
- PERF_CTL[2]: 7 multi-event per clock maximum.
- PERF_CTL[3]: 63 multi-event per clock maximum.
- PERF_CTL[4]: 7 multi-event per clock maximum.
- PERF_CTL[5]: 7 multi-event per clock maximum.

Not all performance monitor events can be counted on all counters. The performance counter registers are generally assigned to specific blocks of the core according to Table 10; however, there are exceptions when an events is implemented by another block of the core and therefore has the counter restrictions of that block. Each core event description starts with one of the following terms to indicate which counters support that event. Selecting an event for a counter that does not support that counter will produce undefined results.

Table 10: Core PMC mapping to PERF_CTL[5:0]

Term	Definition
PERF_CTL[5:0]	PERF_CTL[5:0] are used to count events in the LS/DC and EX where the number of events logged per cycle can vary up to 7.
PERF_CTL[3,0]	PERF_CTL[3,0] are used to count events in the LS/DC and EX where the number of events logged per cycle can vary up to 31.
PERF_CTL[0]	PERF_CTL[0] are used to count events in the LS/DC, EX, IF/DE and CU where the number of events logged per cycle can vary up to 31.

Table 10: Core PMC mapping to PERF_CTL[5:0]

Term	Definition
PERF_CTL[3]	PERF_CTL[3] are used to count events in the LS/DC, EX and FP where the number of events logged per cycle can vary up to 63.
PERF_CTL[2:0]	PERF_CTL[2:0] are used to count events in the IF/DE and CU; The number of events logged per cycle can vary up to 7.
PERF_CTL[5:3]	PERF_CTL[5:3] are used to count events in the FP; The number of events logged per cycle can vary up to 7.

Writing the performance counters can be useful if there is an intention for software to count a specific number of events, and then trigger an interrupt when that count is reached. An interrupt can be triggered when a performance counter overflows. Software should use the WRMSR instruction to load the count as a two's-complement negative number into the performance counter. This causes the counter to overflow after counting the appropriate number of times.

In addition to the RDMSR instruction, the PERF_CTR[5:0] registers can be read using a special read performance-monitoring counter instruction, RDPMC.

The accuracy of the performance counters is not ensured. The performance counters are not assured of producing identical measurements each time they are used to measure a particular instruction sequence, and they should not be used to take measurements of very small instruction sequences. The RDPMC instruction is not serializing, and it can be executed out-of-order with respect to other instructions around it. Even when bound by serializing instructions, the system environment at the time the instruction is executed can cause events to be counted before the counter value is loaded into EDX:EAX.

2.6.2 NB Performance Monitor Counters

The NB performance monitor counters are used by software to count specific events that occur in the NB. Each node provides four 48-bit performance counters.

[MSRC001_024\[6,4,2,0\]](#) [Northbridge Performance Event Select (NB_PERF_CTL[3:0])] and [MSRC001_024\[7,5,3,1\]](#) [Northbridge Performance Event Counter (NB_PERF_CTR[3:0])] specify the events to be monitored and how they are monitored.

All of the events are specified in [3.23 \[NB Performance Counter Events\]](#).

All NB performance monitor events can be counted on all counters. E.g. Unlike core performance events, there are no NB events that can not be counted on some NB performance counters.

All NB performance events are one event per clock. E.g. Unlike core performance events, there are no NB multi-event performance events.

NB performance counters do not support APIC interrupt capability.

In addition to the RDMSR instruction, the NB_PERF_CTR[3:0] registers can be read using a special read performance-monitoring counter instruction, RDPMC.

The accuracy of the performance counters is not ensured. The performance counters are not assured of producing identical measurements each time they are used to measure a particular instruction sequence, and they should not be used to take measurements of very small instruction sequences. The RDPMC instruction is not

serializing, and it can be executed out-of-order with respect to other instructions around it. Even when bound by serializing instructions, the system environment at the time the instruction is executed can cause events to be counted before the counter value is loaded into EDX:EAX.

2.6.3 Instruction Based Sampling (IBS)

IBS is a code profiling mechanism that enables the processor to select a random instruction fetch or micro-op after a programmed time interval has expired and record specific performance information about the operation. An interrupt is generated when the operation is complete as specified by [MSRC001_103A \[IBS Control\]](#). An interrupt handler can then read the performance information that was logged for the operation.

The IBS mechanism is split into two parts: instruction fetch performance controlled by [MSRC001_1030 \[IBS Fetch Control \(IC_IBS_CTL\)\]](#); and instruction execution performance controlled by [MSRC001_1033 \[IBS Execution Control \(SC_IBS_CTL\)\]](#). Instruction fetch sampling provides information about instruction TLB and instruction cache behavior for fetched instructions. Instruction execution sampling provides information about micro-op execution behavior. The data collected for instruction fetch performance is independent from the data collected for instruction execution performance. Support for the IBS feature is indicated by the [CPUID Fn8000_0001_ECX\[IBS\]](#).

Instruction fetch performance is profiled by recording the following performance information for the tagged instruction fetch:

- If the instruction fetch completed or was aborted. See [MSRC001_1030](#).
- The number of clock cycles spent on the instruction fetch. See [MSRC001_1030](#).
- If the instruction fetch hit or missed the IC, hit/missed in the L1 and L2 TLBs, and page size. See [MSRC001_1030](#).
- The linear address, physical address associated with the fetch. See [MSRC001_1031](#), [MSRC001_1032](#).

Instruction execution performance is profiled by tagging one micro-op associated with an instruction. Instructions that decode to more than one micro-op return different performance data depending upon which micro-op associated with the instruction is tagged. These micro-ops are associated with the RIP of the next instruction to retire. The following performance information is returned for the tagged micro-op:

- Branch and execution status for micro-ops. See [MSRC001_1035](#).
- Branch target address for branch micro-ops. See [MSRC001_103B](#).
- The logical address associated with the micro-op. See [MSRC001_1034](#).
- The linear and physical address associated with a load or store micro-op. See [MSRC001_1038](#), [MSRC001_1039](#).
- The data cache access status associated with the micro-op: DC hit/miss, DC miss latency, TLB hit/miss, TLB page size. See [MSRC001_1037](#).
- The number clocks from when the micro-op was tagged until the micro-op retires. See [MSRC001_1035](#).
- The number clocks from when the micro-op completes execution until the micro-op retires. See [MSRC001_1035](#).
- Source information for DRAM, MMIO, L3 hit and state. See [MSRC001_1036](#).

2.7 Configuration Space

PCI-defined configuration space was originally defined to allow up to 256 bytes of register space for each function of each device; these first 256 bytes are called base configuration space (BCS). It was expanded to support up to 4096 bytes per function; bytes 256 through 4095 are called extended configuration space (ECS). The processor includes configuration space registers located in both BCS and ECS. Processor configuration space is accessed through bus 0, devices 18h to 1Fh, where device 18h corresponds to node 0 and device 1Fh corresponds to node 7. See [2.7.3 \[Processor Configuration Space\]](#).

Configuration space is accessed by the processor through two methods:

- IO-space configuration: IO instructions to addresses CF8h and CFCh.
 - Enabled through [IOCF8 \[IO-Space Configuration Address\]\[ConfigEn\]](#), which allows access to BCS.
 - Access to ECS enabled through [MSRC001_001F \[Northbridge Configuration 1 \(NB_CFG1\)\]\[EnableCf8ExtCfg\]](#).
 - Use of IO-space configuration can be programmed to generate GP faults through [MSRC001_0015 \[Hardware Configuration \(HWCR\)\]\[IoCfgGpFault\]](#).
 - SMI trapping for these accesses is specified by [MSRC001_0054 \[IO Trap Control \(SMI_ON_IO_TRAP_CTL_STS\)\]](#) and [MSRC001_00\[53:50\] \[IO Trap \(SMI_ON_IO_TRAP_\[3:0\]\)\]](#).
- MMIO configuration: configuration space is a region of memory space.
 - The base address and size of this range is specified by [MSRC001_0058 \[MMIO Configuration Base Address\]](#). The size is controlled by the number of configuration-space bus numbers supported by the system. Accesses to this range are converted configuration space as follows:
 - Address[31:0] = {0h, bus[7:0], device[4:0], function[2:0], offset[11:0]}.

BIOS may use either configuration space access mechanism during boot. Before booting the OS, BIOS must disable IO access to ECS, enable MMIO configuration and build an ACPI defined MCFG table. BIOS ACPI code must use MMIO to access configuration space.

Per the link specification, BCS accesses utilize link addresses starting at FD_FE00_0000h and ECS accesses utilize link addresses starting at FE_0000_0000h.

2.7.1 MMIO Configuration Coding Requirements

MMIO configuration space accesses must use the uncacheable (UC) memory type. Instructions used to read MMIO configuration space are required to take the following form:

```
mov eax/ax/al, <any_address_mode>;
```

Instructions used to write MMIO configuration space are required to take the following form:

```
mov <any_address_mode>, eax/ax/al;
```

No other source/target registers may be used other than eax/ax/al.

In addition, all such accesses are required not to cross any naturally aligned DW boundary. Access to MMIO configuration space registers that do not meet these requirements result in undefined behavior.

2.7.2 MMIO Configuration Ordering

Since MMIO configuration cycles are not serializing in the way that IO configuration cycles are, their ordering rules relative to posted may result in unexpected behavior.

Therefore, processor MMIO configuration space is designed to match the following ordering relationship that exists naturally with IO-space configuration: if a core generates a configuration cycle followed by a posted-write cycle, then the posted write is held in the processor until the configuration cycle completes. As a result, any unexpected behavior that might have resulted if the posted-write cycle were to pass MMIO configuration cycle is avoided.

2.7.3 Processor Configuration Space

The processor includes configuration space as described in 3 [Registers]. Accesses to unimplemented registers of implemented functions are ignored: writes dropped; reads return 0's. Accesses to unimplemented functions also ignored: writes are dropped; however, reads return all F's. The processor does not log any master abort events for accesses to unimplemented registers or functions.

Accesses to device numbers of devices not implemented in the processor are routed based on the configuration map registers. If such requests are master aborted, then the processor can log the event.

2.8 Northbridge (NB)

Each node includes a single northbridge that provides the interface to the local core(s), the interface to system memory, and the interface to system IO devices. The NB includes all power planes except VDD; see [2.5.1 \[Processor Power Planes And Voltage Control\]](#).

The NB is responsible for routing transactions sourced from cores and link to the appropriate core, cache, DRAM, or link. See [2.4.5 \[System Address Map\]](#).

2.8.1 NB Architecture

Major NB blocks are: System Request Interface (SRI), Memory Controller (MCT), DRAM Controllers (DCTs), and crossbar (XBAR). SRI interfaces with the core(s). MCT maintains cache coherency and interfaces with the DCTs; MCT maintains a queue of incoming requests called MCQ. XBAR is a switch that routes packets between SRI, MCT, and the link.

The MCT operates on physical addresses. Before passing transactions to the DCTs, the MCT converts physical addresses into *normalized* addresses that correspond to the values programmed into [D18F2x\[5C:40\]_dct\[1:0\] \[DRAM CS Base Address\]](#). Normalized addresses include only address bits within the DCTs' range. The normalized address varies based on DCT interleave and hoisting settings in [D18F2x110 \[DRAM Controller Select Low\]](#) and [D18F2x114 \[DRAM Controller Select High\]](#) as well as node interleaving based on [D18F1x\[144:140,44:40\] \[DRAM Base/Limit\]](#).

2.8.2 NB Routing

2.8.2.1 Address Space Routing

There are four main types of address space routed by the NB:

1. Memory space targeting system DRAM
2. Memory space targeting IO (MMIO)
3. IO space
4. Configuration space.

2.8.2.1.1 DRAM and MMIO Memory Space

Memory space transactions provide the NB with the physical address, cacheability type, access type, and DRAM/MMIO destination type as specified in section [2.4.5.1.2 \[Determining The Access Destination for Core Accesses\]](#).

Memory-space transactions are handled by the NB as follows:

- IO-device accesses are compared against:
 - If the access matches [D18F1x\[1CC:180,BC:80\] \[MMIO Base/Limit\]](#), then the transaction is routed to the root complex;
 - Else, if the access matches [D18F1x\[144:140,44:40\] \[DRAM Base/Limit\]](#), then the access is routed to the DCT;
 - Else, the access is routed to the UMI.
- For core accesses the routing is determined based on the DRAM/MMIO destination:
 - If the destination is DRAM:
 - If the access matches [D18F1x\[144:140,44:40\] \[DRAM Base/Limit\]](#), then the transaction is routed to the DCT;
 - Else, the access is routed to the UMI.
 - If the destination is MMIO:
 - If the access matches [D18F1x\[1CC:180,BC:80\] \[MMIO Base/Limit\]](#), then the transaction is routed

to the root complex;

- Else, If the access matches the VGA-compatible MMIO address space and **D18F1xF4**[VE]=1 then **D18F1xF4** describes how the access is routed and controlled.
- Else, the access is routed to the UMI.

2.8.2.1.2 IO Space

IO-space transactions from IO links or cores are routed as follows:

- If the access matches **D18F1x[DC:C0]** [IO-Space Base/Limit], then the transaction is routed to the root complex;
- Else, If the access matches the VGA-compatible IO address space and **D18F1xF4**[VE]=1 then **D18F1xF4** describes how the access is routed and controlled.
- Else, the access is routed to the UMI.

2.8.2.1.3 Configuration Space

Configuration-space transactions from IO links are master aborted. Configuration-space transactions from cores are routed as follows:

- If the access targets the configuration space of an existing node (based on the configuration-space address and **D18F0x60**[NodeCnt]), then it is routed to that node.
- Else, if the access matches **D18F1x[EC:E0]** [Configuration Map], then the transaction is routed to the specified link;
- Else, the access is routed to the node or link that contains compatibility (subtractive) address space, specified by **D18F0x60**[SbNode] and **D18F0x64**[SbLink].

2.9 DRAM Controllers (DCTs)

The processor includes two DRAM controllers (DCTs). Each DCT controls one 64-bit DDR3 DRAM channel. A DRAM channel consists of the group of DRAM interface pins connecting to one series of DIMMs.

The DCTs operate on normalized addresses corresponding to the values programmed into the [D18F2x\[5C:40\]_dct\[1:0\] \[DRAM CS Base Address\]](#). Normalized addresses only include address bits within a DCT's range. The physical to normalized address translation varies based on DCT interleave and hoisting settings. See [2.9.6 \[Memory Interleaving Modes\]](#) and [2.9.7 \[Memory Hoisting\]](#).

The following restrictions limit the DIMM types and configurations supported by the DCTs:

- All DIMMs connected to the processor are required to operate at the same MEMCLK frequency, regardless of the channel. All DCTs must be programmed to the same frequency.

Table 11: DCT Definitions

Term	Definition
SR	Single Rank
DR	Dual Rank
any	Any: SR or DR .
MRS	JEDEC defined DRAM Mode Register Set.
NP	No DIMM populated
DIMM0	DIMM slots 0-n. The DIMMs on each channel are numbered from 0 to n where DIMM0 is the DIMM closest to the processor on that channel and DIMMn is the DIMM farthest from the processor on that channel.
DIMM1	
DimmsPopulated	The number of DIMMs populated per channel
NumDimmSlots	The number of motherboard DIMM slots per channel
DIMM	The DIMM being configured
Rank	The rank being configured
NumRegisters	The number of registers on the DIMM being configured
NumDimmSlots	The number of motherboard DIMM slots per channel
RowAddrBits	SPDByte[5][5:3] of the DIMM being configured.
DramCapacity	SPDByte[4][3:0] of the DIMM being configured.
RankMap	SPDByte[63][0] of the DIMM being configured.
NumRanks	SPDByte[7][5:3] of the DIMM being configured, or the number of ranks soldered down.
DeviceWidth	SPDByte[7][2:0] of the DIMM being configured.
AutoSelfRefresh	SPDByte[31][2] of the DIMM being configured.
ExtendedTemperature-Range	SPDByte[31][0] of the DIMM being configured.
DdrRate	The DDR data rate (MT/s).
VDDIO	DDR VDDIO in V.
UDIMM	DCT is configured for UDIMM if (D18F2x90_dct[1:0][UnbuffDimm] ==1).
SODIMM	DCT is configured for UDIMM if (D18F2x90_dct[1:0][UnbuffDimm] ==1) and the processor is connected to SO-DIMM(s) . This definition is used to distinguish between settings optimized for SO-DIMMs, as in a laptop, or standard sized UDIMMs as in a desktop.

The following restrictions limit the DIMM types and configurations supported by the DCTs:

- All DIMMs connected to a processor are required to operate at the same MEMCLK frequency, regardless of the channel. All DCTs must be programmed to the same frequency.

The tables below list the maximum DIMM speeds supported by the processor for different configurations. The motherboard should comply with the relevant AMD socket motherboard design guideline (MBDG) to achieve the rated speeds. In cases where MBDG design options exist, lower-quality options may compromise the maximum achievable speed; motherboard designers should assess the tradeoffs.

Table 12: DDR3 unbuffered DIMM maximum frequency support for FM2 (per channel)

DIMM Slots	Installed DIMMs	Frequency (MT/s) ⁵		
		1.5V	1.35V	1.25V
2	1	2133	1600	1333
	2 SR	1866	1600	1333
	1 DR + 1 SR/DR	(TN-A0 TN-A1) ⁶ : 1600 RL-A1 ⁶ : 1866	1600	1066
5. Population restrictions (including the order for partially populated channels) may apply. See 2.9.5.6.6.				
6. See 1.5 for revision information.				

Table 13: DDR3 SO-DIMM maximum frequency support for FS1r2 (per channel)

DIMM Slots	Installed DIMMs	Frequency (MT/s) ⁵		
		1.5V	1.35V	1.25V
1	1	1866	1600	1333
7. Population restrictions (including the order for partially populated channels) may apply. See 2.9.5.6.6.				

Table 14: DDR3 unbuffered DIMM maximum frequency support for FP2 (per channel) for (TN-A0 || TN-A1)

DIMM Slots	Installed DIMMs	Frequency (MT/s) ⁸		
		1.5V	1.35V	1.25V
1	1	1600	1600	1333
2	1	1333	1333	1333
	2 SR	1333	1333	1333
	1 DR + 1 SR/DR	1333	1333	1066
8. For (TN-A0 TN-A1). See 1.5 for revision information. Population restrictions (including the order for partially populated channels) may apply. See 2.9.5.6.6.				

Table 15: DDR3 SO-DIMM maximum frequency support for FP2 (per channel) for (TN-A0 || TN-A1)

DIMM Slots	Installed DIMMs	Frequency (MT/s) ⁹		
		1.5V	1.35V	1.25V
1	1	1600	1600	1333
2	1	1333	1333	1066
	2 SR	1333	1333	1066
	1 DR + 1 SR/DR	1333	1333	1066

9. For (TN-A0 || TN-A1). See 1.5 for revision information. Population restrictions (including the order for partially populated channels) may apply. See 2.9.5.6.6.

Table 16: DDR3 unbuffered DIMM maximum frequency support for FP2 (per channel) for RL-A1

DIMM Slots	Installed DIMMs	Frequency (MT/s) ⁵		
		1.5V	1.35V	1.25V
1	1	1600	1600	1333
2	1	1600	1333	1333
2	2 SR	1600	1333	1333
	1 DR + 1 SR/DR	1600	1333	1066

1. For RL-A1. See 1.5 for revision information. Population restrictions (including the order for partially populated channels) may apply. See 2.9.5.6.6.

Table 17: DDR3 SO-DIMM maximum frequency support for FP2 (per channel) for RL-A1

DIMM Slots	Installed DIMMs	Frequency (MT/s) ⁵		
		1.5V	1.35V	1.25V
1	1	1600	1600	1333
2	1	1333	1333	1333
2	2 SR	1333	1333	1333
	1 DR + 1 SR/DR	1333	1333	1066

1. For RL-A1. See 1.5 for revision information. Population restrictions (including the order for partially populated channels) may apply. See 2.9.5.6.6.

Table 18: DDR3 Soldered-down DRAM maximum frequency support for FP2 (per channel)

Ranks	Frequency (MT/s) ⁵		
	1.5V	1.35V	1.25V
Up to 4	1333	1333	1333

1. Population restrictions (including the order for partially populated channels) may apply. See 2.9.5.6.6.

2.9.1 DCT Configuration Registers

There are two types of DCT configuration registers:

- Registers for which there is one instance for all DCT's. E.g. [D18F2xAC \[DRAM Controller Temperature Status\]](#).
- Registers for which there is one instance per DCT. E.g. [D18F2x78_dct\[1:0\] \[DRAM Control\]](#).
 - For [D18F2x78_dct\[x\]](#), $x = \text{D18F1x10C}[\text{DctCfgSel}]$; see [D18F1x10C\[DctCfgSel\]](#).
 - The syntax for this register type is described by example as follows:
 - [D18F2x78_dct\[1:0\]](#) refers to all instances of the [D18F2x78](#) register.
 - [D18F2x78_dct\[1\]](#) refers to the [D18F2x78](#) register instance for DCT1.

Registers for which there is one instance per memory P-state use [D18F1x10C\[MemPsSel\]](#) for software accesses. The syntax for this register type is described by example as follows:

- [D18F2x2E8_dct\[1:0\]_mp\[1:0\]](#) refers to all instances of the [D18F2x2E8](#) register.
- [D18F2x2E8_dct\[1:0\]_mp\[1\]](#) refers to the register for memory P-state 1 of either or both DCTs.
- Software must ensure consistency of context when accessing registers with both controller fields and DDR phy fields. See [D18F2x98_dct\[1:0\]](#) for more information.

2.9.2 DDR Pad to Processor Pin Mapping

The relationship of pad drivers to processor pins varies by package as shown in the following table.

Table 19: Package pin mapping

Pad	Pin ¹		
	FS1r2	FM2	FP2
MEMCLK0_H[0]	MA_CLK_H[0]		
MEMCLK0_H[1]	MA_CLK_H[1]		
MEMCLK0_H[2]	NC	MA_CLK_H[2]	
MEMCLK0_H[3]	NC	MA_CLK_H[3]	
MEMCLK0_H[4]	NC		
MEMCLK0_H[5]	NC		
MEMCLK0_H[6]	NC		
MEMCLK0_H[7]	NC		
MEMCLK1_H[0]	MB_CLK_H[0]		
MEMCLK1_H[1]	MB_CLK_H[1]		
MEMCLK1_H[2]	NC	MB_CLK_H[2]	
MEMCLK1_H[3]	NC	MB_CLK_H[3]	
MEMCLK1_H[4]	NC		
MEMCLK1_H[5]	NC		
MEMCLK1_H[6]	NC		
MEMCLK1_H[7]	NC		
MEMCS0_L[0]	MA0_CS_L[0]		
MEMCS0_L[1]	MA0_CS_L[1]		
MEMCS0_L[2]	NC	MA1_CS_L[0]	
MEMCS0_L[3]	NC	MA1_CS_L[1]	
MEMCS0_L[4]	NC		

Table 19: Package pin mapping

Pad	Pin ¹		
	FS1r2	FM2	FP2
MEMCS0_L[5]	NC		
MEMCS0_L[6]	NC		
MEMCS0_L[7]	NC		
MEMCS1_L[0]	MB0_CS_L[0]		
MEMCS1_L[1]	MB0_CS_L[1]		
MEMCS1_L[2]	NC	MB1_CS_L[0]	
MEMCS1_L[3]	NC	MB1_CS_L[1]	
MEMCS1_L[4]	NC		
MEMCS1_L[5]	NC		
MEMCS1_L[6]	NC		
MEMCS1_L[7]	NC		
MEMODT0[0]	MA0_ODT[0]		
MEMODT0[1]	MA0_ODT[1]		
MEMODT0[2]	NC	MA1_ODT[0]	
MEMODT0[3]	NC	MA1_ODT[1]	
MEMODT1[0]	MB0_ODT[0]		
MEMODT1[1]	MB0_ODT[1]		
MEMODT1[2]	NC	MB1_ODT[0]	
MEMODT1[3]	NC	MB1_ODT[1]	
MEMCKE0[0]	MA_CKE[0]		
MEMCKE0[1]	MA_CKE[1]		
MEMCKE0[2]	NC		
MEMCKE0[3]	NC		
MEMCKE1[0]	MB_CKE[0]		
MEMCKE1[1]	MB_CKE[1]		
MEMCKE1[2]	NC		
MEMCKE1[3]	NC		

1. For differential pins, only positive polarity pins are shown; negative polarity pins have corresponding mapping and are controlled by the same CSR field.
NC = Not connected. BIOS should tri-state or disable the pad for maximum power savings.

2.9.2.1 DDR Chip to Pad Mapping

The relationship of chip to pad drivers is shown in the following table. BIOS should disable or power down unused chips for maximum power savings.

Table 20: Pad (from chiplet) pin mapping

Chip	Group	Pad ¹		
		FS1r2	FM2	FP2
Clk chip 0	0	MEMCLK0_H[1:0]		
Clk chip 1	0	MEMCLK0_H[3:2]		
Cmd/Addr 0	0	MEMCS0_L[7,5,3,1]		
	1	MEMODT0[3:0]		
	2	MEMCS0_L[6,4,2,0]		
Cmd/Addr 1	0	MEMRAS0_L		
		MEMCAS0_L		
		MEMWE0_L		
		MEMADD0[13]		
	1	MEMADD0[10,0]		
		MEMBANK0[1:0]		
2	MEMPAR0			
Address 2	0	MEMCKE0[3:0]		
	1	MEMADD0[0,1,8,9]		
	2	MEMADD0[2,3,10,11]		
	3	MEMADD0[4,5,12,13]		
	4	MEMADD0[6,7,14,15]		
Data ²	0	MEMDATA[5,4,1,0]		
	1	MEMDQS_H[0], MEMDQSDM[0]		
	2	MEMDATA[7,6,3,2]		
<p>1. For differential pads, only positive polarity pads are shown; negative polarity pads have corresponding mapping and are controlled by the same chip. Only channel A map is shown. Channel B is similar.</p> <p>2. Only Data chip 0 is shown. Data chips [7:1] are repeated with sequential DQ/DQS/DM pin numbers.</p>				

2.9.3 DRAM Controller Direct Response Mode

The DCT supports direct response mode for responding to a cache line fill request before the DCT is initialized. In direct response mode, the target DCT responds to a cache line fill request by returning 64 bytes of all ones without issuing a read transaction on the DRAM bus. The BIOS uses this feature to allocate cache lines for temporary data storage. The controller exits direct response mode when either `D18F2x7C_dct[1:0][EnDramInit]` or `D18F2x90_dct[1:0][InitDram]` is set to 1.

See [2.9.5.8 \[DRAM Device and Controller Initialization\]](#) and [2.3.3 \[Using L2 Cache as General Storage During Boot\]](#).

2.9.4 DRAM Data Burst Mapping

DRAM requests are mapped to data bursts on the DDR bus in the following order:

- When `D18F2x110[DctDatIntLv] = 0`, a 64 B request is mapped to each of the eight sequential data beats as QW0, QW1...QW7.

- When `D18F2x110[DctDatIntLv] = 1`, the order of cache data to QW on the bus is the same except that even and odd bits are interleaved on the DRAM bus as follows:
 - For every 8 bytes in the cache line, even bits map to QW0, QW2, QW4, and QW6 on the DRAM bus.
 - For every 8 bytes in the cache line, odd bits map to QW1, QW3, QW5, and QW7 on the DRAM bus.

2.9.5 DCT/DRAM Initialization and Resume

DRAM initialization involves several steps in order to configure the DRAM controllers and the DRAM, and to tune the DRAM channel for optimal performance. DRAM resume requires several steps to configure the DCTs to properly resume from the S3 state. The following sequence describes the steps needed after a reset for initialization or resume:

- To disable an unused DRAM channel see [DRAM Channel Disable \[2.9.5.10\]](#).
1. Configure the DDR supply voltage regulator. See [2.9.5.1](#).
 2. Force NB P-state to NBP0. See [2.9.5.2](#).
 3. Force register accesses to M0. See [2.9.5.3](#) for further requirements used in steps below.
 4. DDR phy initialization. See [2.9.5.4](#).
 5. DRAM device and controller initialization.
 - If BIOS is booting from an unpowered state (ACPI S4, S5 or G3), then it performs the following:
 - a. Program SPD configuration. See [2.9.5.5](#).
 - b. Program Non-SPD configuration. See [2.9.5.6](#).
 - c. Program DCT training specific configuration. See [2.9.5.7](#).
 - d. Program the remaining DCT registers not covered by an explicit sequence dependency.
 - e. DRAM device initialization. See [2.9.5.8](#).
 - If BIOS is resuming the platform from S3 state, then it performs the following:
 - a. Restore all DCT and phy registers that were programmed during the first boot from non-volatile storage. See [2.9.5.5](#) and [2.9.5.6](#) for a review of registers.
 - b. Program `D18F2x90_dct[1:0][ExitSelfRef] = 1`.
 - c. Restore the trained delayed values (found during the initial boot in step 6 below) from nonvolatile storage.
 - d. Continue at step 9.
 6. DRAM data training.
 - A. Write levelization training. See [2.9.5.9.1](#).
 - B. DQS receiver enable training. See [2.9.5.9.2](#).
 - C. DQS position training. See [2.9.5.9.4](#).
 - D. MaxRdLatency training. See [2.9.5.9.5.1](#).
 7. NB P-state specific training. For each NB P-state from NBP1 to `D18F5x170[NbPstateMaxVal]`:
 - A. Force the NB P-state. See [2.9.5.2](#).
 - B. MaxRdLatency training. See [2.9.5.9.5.1](#).
 8. Release NB P-state force.
 9. Program DCT for normal operation. See [2.9.5.7](#).
 10. Program DRAM phy for power savings. See [2.9.5.11](#).

The DRAM subsystem is ready for use.

2.9.5.1 Low Voltage DDR3

The processor supports JEDEC defined DDR3L and DDR3U devices which operate at voltages lower than 1.5V.

Platforms supporting low voltage devices should power up VDDIO at 1.35V. BIOS should not operate DIMMs at voltages higher than supported as indicated by SPD Byte 6: Module Nominal Voltage, VDD.

BIOS should consult vendor data sheets for the supply voltage regulator programming requirements. On supported platforms, BIOS must take steps to configure the supply voltage regulators as follows:

1. Read the SPD of all DIMMs within the programmable VDDIO domain and check all of the defined bits within the SPD byte to determine the common operating voltages.
2. Configure VDDIO to match the lowest common supported voltage based on the SPD values.
 - If the DIMMs do not specify a common operating voltage then BIOS must take platform vendor defined action to notify the end user of the mismatch and to protect DIMMs from damage.
3. Additional derating of the DDR speed may be necessary for reliable operation at lower voltage. See [2.9.5.6.6 \[DRAM Address Timing and Output Driver Compensation Control\]](#) for information on voltage and electrical load specific maximum speed.

2.9.5.2 NB P-state Specific Configuration

A subset of DCT configuration and training must be repeated for each enabled NB P-state. To accomplish this, BIOS forces the processor to the desired NB P-state and releases the force once DRAM initialization and training is complete.

When [D18F5x174\[NbPstateDis\]](#)=0, BIOS performs the following to configure NB P-states for MEMCLK frequency specific operation:

1. Adjust the NB P-state northbridge voltage. See [2.5.4.1.2.2](#).
 - BIOS may perform this step once for all defined NB P-states, or for each time an NB P-state is forced during DRAM initialization to the target state only.

When [D18F5x174\[NbPstateDis\]](#)=0, BIOS performs the following to configure the DCT for NB P-states:

Before DRAM training, BIOS performs the following to transition to the highest performance NB P-state, NBP0:

1. Save the initial value of [D18F5x170](#).
2. If [D18F5x174\[CurNbPstate\]](#)!= NBP0 and [D18F5x174\[CurNbPstate\]](#)==[D18F5x170\[NbPstateHi\]](#), transition to target NB P-state, NBP0, via Steps 3-7 below.
Else if [D18F5x174\[CurNbPstate\]](#)!= NBP0 and [D18F5x174\[CurNbPstate\]](#)==[D18F5x170\[NbPstateLo\]](#), transition to target NB P-state, NBP0, via Steps 3,6-7 below.
3. Program the configuration registers which contain multiple internal copies for each NB P-state. See [D18F1x10C\[NbPsSel\]](#).
4. Program [D18F5x170](#) to transition the NB P-state:
 - [SwNbPstateLoDis](#) = [NbPstateDisOnP0](#) = [NbPstateThreshold](#) = 0.
 - [NbPstateLo](#) = [NbPstateMaxVal](#).
5. Wait for [D18F5x174\[CurNbPstate\]](#) to equal [NbPstateLo](#).
6. Program [D18F5x170](#) to force the NB P-state:
 - [NbPstateHi](#) = target NB P-state.
 - [SwNbPstateLoDis](#) = 1
7. Wait for [D18F5x174\[CurNbPstate\]](#) to equal the target NB P-state.

BIOS performs the following to release the NB P-state force:

8. Restore the initial [D18F5x170\[SwNbPstateLoDis, NbPstateDisOnP0, NbPstateLo\]](#) values.
9. Restore the initial [D18F5x170\[NbPstateThreshold, NbPstateHi\]](#) values.

See also 2.5.4.1 [NB P-states].

2.9.5.3 Memory P-state Specific Configuration

A subset of DCT configuration and training must be repeated for each enabled memory P-state. See also 2.5.7.1 [Memory P-states]. To accomplish this, BIOS forces the processor to the NBP0 state and uses the M0 context to train for values that are used by all memory P-states. For example, BIOS uses D18F2x94_dct[1:0][MemClkFreq] to specify the dram frequency during each pass of training. BIOS saves the trained configuration values at each step and uses them appropriately as outlined below:

1. Force NB P-state to NBP0. See 2.9.5.2.
2. Program D18F1x10C[MemPsSel]=0. See 2.9.1.
3. Program D18F2x9C_x0D04_E008_dct[1:0][PStateToAccess]=0.
4. Using the M0 context, BIOS optimally configures the controller and trains at DDR667 according to 2.9.5 steps 4 thru 7.
 - A. BIOS saves the trained values for use in M1 and for subsequent higher frequencies.
 - B. BIOS saves the controller configuration for use in M1.
 - C. BIOS optimally writes all the values to registers in step 6. It may optionally write the values as each is computed or trained. If this is done then BIOS must ensure that D18F1x10C[MemPsSel] and D18F2x9C_x0D04_E008_dct[1:0][PStateToAccess] are configured correctly before and after each register access.
5. BIOS increases the frequency and repeats step 4 until the target frequency is trained. The target frequency is always associated with M0.
6. BIOS programs the M1 controller context with the DDR667 configuration and trained results. This must be done before the next subsequent NB P-state change.
 - a. Program D18F1x10C[MemPsSel]=1.
 - b. Program D18F2x9C_x0D04_E008_dct[1:0][PStateToAccess] = 1.
 - c. Program D18F2x2E0_dct[1:0][M1MemClkFreq] and D18F2x9C_x0D0F_E000_dct[1:0]_mp[1][Rate, FreqValid] = {04h, 1}.
 - d. Program DCT configuration and trained delay values saved earlier. Three groups of registers must be written in the same order as the initialization sequence:
 - Fence values: D18F2x9C_x0000_000C_dct[1:0], D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0], D18F2x9C_x0D0F_E019_dct[1:0].
 - Address/Command timing delays: D18F2x9C_x0000_0004_dct[1:0]_mp[1:0].
 - All other Data/Dqs delays: See step 6 in section 2.9.5.
 - e. Program D18F2x9C_x0D04_E008_dct[1:0][PStateToAccess] = 0.
 - f. Program D18F1x10C[MemPsSel] = 0.

2.9.5.4 DDR Phy Initialization

The BIOS initializes the phy and the internal interface from the DCT to the phy, including the PLLs and the fence value, after each reset and for each time a MEMCLK frequency change is made.

BIOS obtains size, loading, and frequency information about the DIMMs and channels using SPDs prior to phy initialization. BIOS then performs the following actions:

1. Program D18F2x9C_x0D0F_E013_dct[1:0] = 0118h.
2. For each byte lane and each memory P-state: Program D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0][RxSsbMntClkEn] = 0.
3. Program D18F2xA8_dct[1:0][MemPhyPlIPdMode] = 00b.
4. Force the phy to M0 with the following sequence:

- A. Program `D18F2x9C_x0D0F_E006_dct[1:0][PIILockTime]` = 190h.
 - B. For each DCT: Program `D18F2x9C_x0000_000B_dct[1:0]` = 80800000h.
 - C. Program `D18F2x9C_x0D0F_E018_dct[0][PhyPSMasterChannel]` = 0.
 - D. Program `D18F2x9C_x0000_000B_dct[0]` = 40000000h.
 - E. For each DCT: Program `D18F2x9C_x0000_000B_dct[1:0]` = 80000000h.
5. Phy Voltage Level Programming. See 2.9.5.4.1.
 6. DRAM channel frequency change. See 2.9.5.4.2.
 7. Phy fence programming. See 2.9.5.4.3.
 8. Phy compensation initialization. See 2.9.5.4.4.

2.9.5.4.1 Phy Voltage Level Programming

BIOS programs the following according to the desired phy VDDIO voltage level:

- Program `D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0][RxVioLvl]`.
- Program `D18F2x9C_x0D0F_[C,8,2][2:0]1F[RxVioLvl]`.
- Program `D18F2x9C_x0D0F_4009_dct[1:0][CmpVioLvl,ComparatorAdjust]`.

See 2.9.5.1 [Low Voltage DDR3].

2.9.5.4.2 DRAM Channel Frequency Change

The following sequence is used to change the DRAM frequency under all boot conditions, including restoring the DCT state when resuming from the S3 state. BIOS performs the sequence only on DCTs with an enabled interface (`D18F2x94_dct[1:0][DisDramInterface]==0`):

For each DCT:

1. Program `D18F2x9C_x0D0F_E006_dct[1:0][PIILockTime]` = 190h.
2. Program `D18F2x94_dct[1:0][MemClkFreqVal]` = 0.
3. Program `D18F2x94_dct[1:0][MemClkFreq]` to the desired DRAM frequency.
4. Program the following parameters which must be configured prior to setting `MemClkFreqVal`:
 - `D18F2x90_dct[1:0][X4Dimm]`
 - `D18F2x94_dct[1:0][ProcOdtDis]`
 - `D18F2x9C_x0000_0004_dct[1:0]_mp[1:0]`
 - `D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0][ProcOdtAdv]`
 - `D18F2x9C_x0D0F_E00A_dct[1:0][SkewMemClk]`
 - `D18F2x210_dct[1:0]_nbp[3:0][RdPtrInit, DataTxFifoWrDly]` of the current NB P-state for the target MEMCLK frequency. See also 2.9.5.2.

For each DCT:

5. Program `D18F2x94_dct[1:0][MemClkFreqVal]` = 1.
 - Wait for `D18F2x94_dct[1:0][FreqChangeInProg]` = 0.

For each DCT:

6. Program `D18F2x9C_x0D0F_E006_dct[1:0][PIILockTime]` according to Table 21.

BIOS must observe the following requirements:

- BIOS must not change the PLL frequency after DRAM has exited from self-refresh.
- BIOS must not change the PLL frequency after the DRAM training for DDR3 DIMMs is complete.

Table 21: DDR PLL Lock Time

D18F2xA8_dct[1:0][MemPhyPllPdMode]	PllLockTime
00b	0Fh
10b	190h

2.9.5.4.2.1 Requirements for DRAM Frequency Change During Training

During DRAM training, BIOS may be required to change the DRAM(MEMCLK) frequency. The steps below describe what is required to prepare the processor and memory subsystem for the new MEMCLK frequency. It is assumed that the memory subsystem has previously been initialized at the current MEMCLK frequency, and this procedure describes only the steps that must be repeated at the new MEMCLK frequency. See [2.9.5.9.1 \[Write Levelization Training\]](#) and [2.9.5.9.2 \[DQS Receiver Enable Training\]](#).

1. Force the NB P-state. See [2.9.5.2](#).
2. Enter self-refresh:
 - A. Program `D18F2x90_dct[1:0][DisDllShutDownSR] = 1`.
 - B. Program `D18F2x90_dct[1:0][EnterSelfRef] = 1`.
 - C. Wait for `D18F2x90_dct[1:0][EnterSelfRef] = 0`.
3. DRAM channel frequency change. See [2.9.5.4.2](#).
4. Change NCLK frequency to meet NCLK-MEMCLK ratio requirements. See [2.5.3.1.6](#).
5. Phy fence programming. See [2.9.5.4.3](#).
6. Phy compensation initialization. See [2.9.5.4.4](#).
7. Program SPD configuration. See [2.9.5.5](#).
8. Program Non-SPD configuration. See [2.9.5.6](#).
9. Exit self-refresh:
 - A. Program `D18F2x90_dct[1:0][ExitSelfRef] = 1`.
 - B. Wait for `D18F2x90_dct[1:0][ExitSelfRef] = 0`.
 - C. Program `D18F2x90_dct[1:0][DisDllShutDownSR] = 0`.
10. Re-program devices with frequency dependent mode register field values. See [2.9.5.8](#).

2.9.5.4.3 Phy Fence Programming

The DDR phy fence logic is used to adjust the phase relationship between the data fifo and the data going to the pad. After any MEMCLK frequency change and before any memory training, BIOS must perform phy fence training for each channel using the following steps:

1. Program `D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0] = 0000_0000h`.
2. Program `D18F2x9C_x0D0F_E019_dct[1:0] = 0000_0000h`.
3. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][FenceTrSel]=10b`.
4. Program `D18F2x9C_x0000_00[51:50]_dct[1:0]=1313_1313h`.
5. Perform phy fence training. See [2.9.5.4.3.1 \[Phy Fence Training\]](#).
6. Write the calculated fence value to `D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxDll]`.
7. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][FenceTrSel]=01b`.
8. Program `D18F2x9C_x0000_00[51:50]_dct[1:0]=1313_1313h`.
9. Perform phy fence training. See [2.9.5.4.3.1 \[Phy Fence Training\]](#).
10. Write the calculated fence value to `D18F2x9C_x0000_000C_dct[1:0][FenceThresholdRxDll]`.
11. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][FenceTrSel]=11b`.
12. Program `D18F2x9C_x0000_00[51:50]_dct[1:0]=1313_1313h`.
13. Perform phy fence training. See [2.9.5.4.3.1 \[Phy Fence Training\]](#).
14. Write the calculated fence value to `D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxPad]`.

15. Program Fence2 threshold for data as follows:
- A. IF (D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxPad] < 16) THEN
 Fence2_TxPad[4:0] = { 1b, D18F2x9C_x0000_000C_dct[1:0][19:16]}
 ELSE
 Fence2_TxPad[4:0] = 00000b
 ENDIF.
 - B. IF (D18F2x9C_x0000_000C_dct[1:0][FenceThresholdRxDll] < 16) THEN
 Fence2_RxDll[4:0] = { 1b, D18F2x9C_x0000_000C_dct[1:0][24:21]}
 ELSE
 Fence2_RxDll[4:0] = 00000b
 ENDIF.
 - C. IF (D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxDll] < 16) THEN
 Fence2_TxDll[4:0] = { 1b, D18F2x9C_x0000_000C_dct[1:0][29:26]}
 ELSE
 Fence2_TxDll[4:0] = 00000b
 ENDIF.
 - D. Program D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0] = {000000b, Fence2_TxDll[4:0],
 Fence2_TxPad[4:0]}.
 - E. Program D18F2x9C_x0D0F_E019_dct[1:0] = {0b, Fence2_RxDll[4:0], 00000b,
 Fence2_TxPad[4:0]}.
16. Reprogram D18F2x9C_x0000_0004_dct[1:0]_mp[1:0].
17. If motherboard routing requires CS[7:6] to adopt address timings, BIOS performs the following:
- A. Program D18F2x9C_x0000_0004_dct[1:0][CSTimingMux67] = 1.
 - B. Program D18F2x9C_x0D0F_8021_dct[1:0]_mp[1:0]:
 - DiffTimingEn = 1.
 - IF (D18F2x9C_x0000_0004_dct[1:0]_mp[1:0][AddrCmdFineDelay] >=
 D18F2x9C_x0D04_E008_dct[1:0][FenceValue]) THEN Fence = 1 ELSE Fence = 0.
 - Delay = D18F2x9C_x0000_0004_dct[1:0]_mp[1:0][AddrCmdFineDelay].
18. Reprogram D18F2x9C_x0000_0004_dct[1:0]_mp[1:0].

When resuming from S3, BIOS reprograms D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxDll, Fence-ThresholdRxDll, FenceThresholdTxPad], D18F2x9C_x0D0F_E019_dct[1:0], and D18F2x9C_x0D0F_8021_dct[1:0]_mp[1:0] from values stored in non-volatile storage instead of training. BIOS reprograms D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0] and D18F2x9C_x0D0F_E019_dct[1:0] as indicated in the steps above.

BIOS may use D18F2x9C_x0000_000C_dct[1:0], D18F2x9C_x0D0F_4003_dct[1], and D18F2x9C_x0D0F_4003_dct[1] for local storage of fence values during the intermediate time between training and writing the values into non-volatile storage in preparation for S3.

2.9.5.4.3.1 Phy Fence Training

1. Program D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][PhyFenceTrEn]=1.
2. Wait 2000 MEMCLKs.
3. Program D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][PhyFenceTrEn]=0.
4. Read the phase recovery engine registers D18F2x9C_x0000_00[51:50]_dct[1:0].
5. Calculate the average of the fine delay values of all byte lanes and subtract 8. If the result is negative then the fence value is zero.

2.9.5.4.4 Phy Compensation Initialization

Each DDR IO driver has a programmable slew rate controlled by the pre-driver calibration code. The recommended slew rate is a function of the DC drive strength. BIOS initializes the recommended nominal slew rate values as follows:

1. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][DisAutoComp] = 1`.
2. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][DisablePredriverCal]=1`.
3. Program TxPreP/TxPreN for Data and DQS according to [Table 22](#), [Table 23](#), or [Table 24](#).
 - A. Program `D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0]={0000b, TxPreP, TxPreN}`.
 - B. Program `D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0]={1000b, TxPreP, TxPreN}`.
4. Program TxPreP/TxPreN for Cmd/Addr according to [Table 25](#), [Table 26](#), or [Table 27](#).
 - A. Program `D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0]={0000b, TxPreP, TxPreN}`.
 - B. Program `D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0]={1000b, TxPreP, TxPreN}`.
5. Program TxPreP/TxPreN for Clock according to [Table 28](#), [Table 29](#), or [Table 30](#).
 - A. Program `D18F2x9C_x0D0F_2[2:0]02_dct[1:0]={1000b, TxPreP, TxPreN}`.
6. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][DisAutoComp] = 0`.

Table 22: Phy predriver calibration codes for Data/DQS at 1.5V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	100_100b	100_100b
	001b	100_100b	100_100b
	010b	100_100b	100_100b
	011b	100_100b	100_100b
1066 - 1333	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b
1600 - 2133	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b

1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN
 See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][DqsDrvStren]
 ELSE
 See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][DataDrvStren]
 ENDIF.
2. See D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0] and
 D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0].

Table 23: Phy predriver calibration codes for Data/DQS at 1.35V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	111_111b	110_110b
	001b	101_101b	101_101b
	010b	101_101b	101_101b
	011b	100_100b	100_100b
1066 - 1333	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b
1600 - 1866	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b
<p>1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][DqsDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][DataDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0] and D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0].</p>			

Table 24: Phy Predriver Calibration Codes for Data/DQS at 1.25V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	111_111b	110_110b
	001b	110_110b	101_101b
	010b	110_110b	101_101b
	011b	100_100b	100_100b
1066 - 1333	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b

Table 24: Phy Predriver Calibration Codes for Data/DQS at 1.25V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
1600 - 1866	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b
<p>1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][DqsDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][DataDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0] and D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0].</p>			

Table 25: Phy predriver calibration codes for Cmd/Addr at 1.5V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	010_010b	010_010b
	001b	010_010b	010_010b
	010b	010_010b	010_010b
	011b	010_010b	010_010b
1066 - 1333	000b	011_011b	011_011b
	001b	011_011b	011_011b
	010b	011_011b	011_011b
	011b	011_011b	011_011b
1600 - 1866	000b	101_101b	101_101b
	001b	101_101b	101_101b
	010b	101_101b	101_101b
	011b	101_101b	101_101b
<p>1. IF (D18F2x9C_x0D0F_C002) THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CkeDrvStren] ELSEIF (D18F2x9C_x0D0F_800[A,6,2]) THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CsOdtDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][AddrCmdDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0] and D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0].</p>			

Table 26: Phy predriver calibration codes for Cmd/Addr at 1.35V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	010_010b	010_010b
	001b	010_010b	010_010b
	010b	010_010b	010_010b
	011b	010_010b	010_010b
1066 - 1333	000b	100_100b	100_100b
	001b	011_011b	011_011b
	010b	011_011b	011_011b
	011b	011_011b	011_011b
1600 - 1866	000b	101_101b	101_101b
	001b	101_101b	101_101b
	010b	100_100b	100_100b
	011b	100_100b	100_100b
<p>1. IF (D18F2x9C_x0D0F_C002)THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CkeDrvStren] ELSEIF (D18F2x9C_x0D0F_800[A,6,2])THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CsOdtDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][AddrCmdDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0] and D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0].</p>			

Table 27: Phy Predriver Calibration Codes for Cmd/Addr at 1.25V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667- 800	000b	010_010b	010_010b
	001b	010_010b	010_010b
	010b	010_010b	010_010b
	011b	010_010b	010_010b
1066 - 1333	000b	110_110b	101_101b
	001b	100_100b	100_100b
	010b	011_011b	011_011b
	011b	010_010b	010_010b

Table 27: Phy Predriver Calibration Codes for Cmd/Addr at 1.25V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
1600 - 1866	000b	111_111b	110_110b
	001b	110_110b	101_101b
	010b	101_101b	100_100b
	011b	101_101b	100_100b
<p>1. IF (D18F2x9C_x0D0F_C002)THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CkeDrvStren] ELSEIF (D18F2x9C_x0D0F_800[A,6,2])THEN See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CsOdtDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][AddrCmdDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0] and D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0].</p>			

Table 28: Phy predriver calibration codes for Clock at 1.5V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	100_100b	100_100b
	001b	100_100b	100_100b
	010b	100_100b	100_100b
	011b	100_100b	100_100b
1066 - 1333	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	101_101b	101_101b
1600 - 1866	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b
<p>1. See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][ClkDrvStren].</p> <p>2. See D18F2x9C_x0D0F_2[2:0]02_dct[1:0].</p>			

Table 29: Phy predriver calibration codes for Clock at 1.35V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	110_110b	101_101b
	001b	110_110b	101_101b
	010b	100_100b	100_100b
	011b	100_100b	100_100b

Table 29: Phy predriver calibration codes for Clock at 1.35V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
1066 - 1333	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	110_110b	101_101b
1600 - 1866	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	110_110b	101_101b

1. See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][ClkDrvStren].
2. See D18F2x9C_x0D0F_2[2:0]02_dct[1:0].

Table 30: Phy Predriver Calibration Codes for Clock at 1.25V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
667 - 800	000b	110_110b	101_101b
	001b	110_110b	101_101b
	010b	100_100b	100_100b
	011b	100_100b	100_100b
1066 - 1333	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b
1600 - 1866	000b	111_111b	110_110b
	001b	111_111b	110_110b
	010b	111_111b	110_110b
	011b	111_111b	110_110b

1. See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][ClkDrvStren].
2. See D18F2x9C_x0D0F_2[2:0]02_dct[1:0].

2.9.5.5 SPD ROM-Based Configuration

The Serial Presence Detect (SPD) ROM is a non-volatile memory device on the DIMM encoded by the DIMM manufacturer. The description of the SPD is usually provided on a data sheet for the DIMM itself along with data describing the memory devices used. The data describes configuration and speed characteristics of the DIMM and the SDRAM components mounted on the DIMM. The associated data sheet also contains the DIMM byte values that are encoded in the SPD on the DIMM.

BIOS reads the values encoded in the SPD ROM through a system-specific interface. BIOS acquires DIMM configuration information, such as the amount of memory on each DIMM, from the SPD ROM on each DIMM and uses this information to program the DRAM controller registers.

For solder-down DRAM, in the absence of an SPD ROM, BIOS provides the information necessary for DRAM configuration. For convenience, this document may refer to solder-down DRAM as a DIMM, and to DRAM data sheet or Jedec DRAM specifications as SPD ROM based, unless otherwise noted.

The SPD ROM provides values for several DRAM timing parameters that are required by the DCT. In general, BIOS should use the optimal value specified by the SPD ROM. These parameters are:

- D18F2x22C_dct[1:0]_mp[1:0][Twr]: Write recovery time
- D18F2x8C_dct[1:0][Tref]
- D18F2x200_dct[1:0]_mp[1:0][Tras]: Active to precharge time
- D18F2x200_dct[1:0]_mp[1:0][Trp]: Precharge time
- D18F2x200_dct[1:0]_mp[1:0][Tred]: RAS to CAS delay
- D18F2x200_dct[1:0]_mp[1:0][Tcl]: CAS latency
- D18F2x204_dct[1:0]_mp[1:0][Trtp]: Internal read to precharge command delay time
- D18F2x204_dct[1:0]_mp[1:0][FourActWindow]: Four activate window delay time
- D18F2x204_dct[1:0]_mp[1:0][Trrd]: Row active to row active delay
- D18F2x204_dct[1:0]_mp[1:0][Trc]: Active to active/refresh time
- D18F2x208_dct[1:0][Trfc3, Trfc2, Trfc1, Trfc0]: Refresh recovery delay time
- D18F2x20C_dct[1:0]_mp[1:0][Twtr]: Internal write to read command delay time

Optimal cycle time is specified for each DIMM and is used to limit or determine bus frequency. See [2.9.5.8 \[DRAM Device and Controller Initialization\]](#).

The DRAM data sheet or the Jedec DRAM specification provides values for some DRAM timing parameters that are required by the DCT, regardless of whether the DRAMs are solder-down. Most of these parameters have a fixed time component as part of their specification. If memory P-states are supported (See [2.5.7.1 \[Memory P-states\]](#)) and a timing register does not have per memory P-state contexts, BIOS must evaluate the minimum value for each frequency and choose the most pessimistic value for all frequencies to program into the register. These parameters are:

- D18F2x20C_dct[1:0]_mp[1:0][Tcwl]: CAS write latency
- D18F2x220_dct[1:0][Tmrd]: MRS command cycle time
- D18F2x220_dct[1:0][Tmod]: MRS command recovery time
- D18F2x224_dct[1:0][Tzqcs]: Short calibration command recovery time
- D18F2x224_dct[1:0][Tzqoper]: Long calibration command recovery time
- D18F2x248_dct[1:0]_mp[1:0][Txpdll]: Exit precharge power down (with DLL frozen) to command delay.
- D18F2x248_dct[1:0]_mp[1:0][Txp]: Exit precharge power down to command delay.
- D18F2x24C_dct[1:0][Tcksrx]: Clock stable to self refresh exit delay
- D18F2x24C_dct[1:0][Tcksre]: Self refresh entry to clock removal delay
- D18F2x24C_dct[1:0][Tckesr]: Self refresh entry to command delay
- D18F2x24C_dct[1:0][Tpd]: Power down entry to exit delay. BIOS may increase this value above the Jedec minimum for certain power down modes.

2.9.5.6 Non-SPD ROM-Based Configuration

There are several DRAM timing parameters and DCT configurations that need to be programmed for optimal memory performance. These values are not derived from the SPD ROM. Several of these timing parameters are functions of other configuration values. These interdependencies must be considered when programming values into several DCT register timing fields. The factors to consider when specifying a value for a specific non-SPD timing parameter are:

- Mixed or non-mixed DIMMs (x4 with x8).

- Training delay values. See 2.9.5.9 [DRAM Training].
- Read and write latency differences.
- The phy's idle clock requirements on the data bus.
- DDR3 ODT timing requirements.
- NCLK frequency
- MEMCLK frequency

The following sub-sections describe how BIOS programs each non-SPD related timing field to a recommended minimum timing value with respect to the above factors.

The following terms are defined to simplify calculations and are calculated in MEMCLKs:

- Latency Difference (LD) = $D18F2x200_dct[1:0]_{mp}[1:0][Tcl] - D18F2x20C_dct[1:0]_{mp}[1:0][Tcwl]$.
- Read ODT Delay (ROD) = $MAX(0, D18F2x240_dct[1:0]_{mp}[1:0][RdOdtOnDuration] - 6) + MAX(0, D18F2x240_dct[1:0]_{mp}[1:0][RdOdtTrnOnDly] - LD)$.
- Write ODT Delay (WOD) = $MAX(0, D18F2x240_dct[1:0]_{mp}[1:0][WrOdtOnDuration] - 6)$.
- $WrEarly = ABS(D18F2x20C_dct[1:0]_{mp}[1:0][WrDqDqsEarly]) / 2$

2.9.5.6.1 TrdrdSdSc, TrdrdSdDc, and TrdrdDd (Read to Read Timing)

The optimal values for $D18F2x218_dct[1:0]_{mp}[1:0][TrdrdSdSc, TrdrdSdDc, TrdrdDd]$ are platform and configuration specific and should be characterized for best performance. Prior to DRAM training, BIOS should program these parameters to the largest defined value. After DRAM training, BIOS should use the guidelines below to configure the recommended platform generic timing values:

- $TrdrdSdSc$ (in MEMCLKs) = 1.
- $TrdrdSdDc$ (in MEMCLKs) = $MAX(TrdrdSdSc, 3 + (IF (D18F2xA8_dct[1:0][PerRankTimingEn]) THEN CEIL(CDD_{TrdrdSdDc} / 2) ELSE 0 ENDIF))$.
- $TrdrdDd$ (in MEMCLKs) = $MAX(TrdrdSdDc, CEIL(MAX(ROD + 3, CDD_{TrdrdDd} / 2 + 3.5)))$.

The Critical Delay Difference (CDD) is the largest delay difference of the channel.

- Each delay difference is $D18F2x9C_x0000_00[2A:10]_{dct}[1:0]_{mp}[1:0][DqsRcvEnGrossDelay]$ minus $D18F2x9C_x0000_00[2A:10]_{dct}[1:0]_{mp}[1:0][DqsRcvEnGrossDelay]$.
- For $CDD_{TrdrdSdDc}$, the subtraction terms are the delays of different chip selects within the same DIMM within the same byte lane.
- For $CDD_{TrdrdDd}$, the subtraction terms are the delays of different DIMMs within the same byte lane.

2.9.5.6.2 TwrwrSdSc, TwrwrSdDc, TwrwrDd (Write to Write Timing)

The optimal values for $D18F2x214_dct[1:0]_{mp}[1:0][TwrwrSdSc, TwrwrSdDc, TwrwrDd]$ are platform and configuration specific and should be characterized for best performance. Prior to DRAM training, BIOS should program these parameters to the largest defined value. After DRAM training, BIOS should use the guidelines below to configure the recommended platform generic timing values:

- $TwrwrSdSc$ (in MEMCLKs) = 1.
- $TwrwrSdDc$ (in MEMCLKs) = $CEIL(MAX(WOD + 3, CDD_{TwrwrSdDc} / 2 + (IF (D18F2xA8_dct[1:0][PerRankTimingEn]) THEN 3.5 ELSE 3 ENDIF)))$.
- $TwrwrDd$ (in MEMCLKs) = $CEIL(MAX(WOD + 3, CDD_{TwrwrDd} / 2 + 3.5))$.

The Critical Delay Difference (CDD) is the largest delay difference of the channel.

- Each delay difference is $D18F2x9C_x0000_00[4A:30]_{dct}[1:0]_{mp}[1:0][WrDqsGrossDly]$ minus $D18F2x9C_x0000_00[4A:30]_{dct}[1:0]_{mp}[1:0][WrDqsGrossDly]$.
- For $CDD_{TwrwrSdDc}$, the subtraction terms are the delays of different chip selects within the same DIMM

within the same byte lane.

- For $CDD_{TwrwrDd}$, the subtraction terms are the delays of different DIMMs within the same byte lane.

BIOS must program these parameters as follows: $TwrwrSdSc \leq TwrwrSdDc \leq TwrwrDd$.

2.9.5.6.3 Twrrd (Write to Read DIMM Termination Turn-around)

The optimal value for $D18F2x218_dct[1:0]_{mp}[1:0][Twrrd]$ is platform and configuration specific and should be characterized for best performance. Prior to DRAM training, BIOS should program these parameters to the largest defined value. After DRAM training, BIOS should use the guidelines below to configure the recommended platform generic timing values:

- $Twrrd$ (in MEMCLKs) = $MAX(1, CEIL(MAX(WOD, CDD_{Twrrd} / 2 + 0.5 - WrEarly) - LD + 3))$.

The Critical Delay Difference (CDD) is the largest delay difference of the channel.

- Each delay difference is $D18F2x9C_x0000_00[4A:30]_dct[1:0]_{mp}[1:0][WrDqsGrossDly]$ minus $D18F2x9C_x0000_00[2A:10]_dct[1:0]_{mp}[1:0][DqsRcvEnGrossDelay]$.
- For CDD_{Twrrd} , the subtraction terms are the delays of different chip selects within the same byte lane.

2.9.5.6.4 TrwtTO (Read-to-Write Turnaround for Data, DQS Contention)

The optimal value for $D18F2x21C_dct[1:0]_{mp}[1:0][TrwtTO]$ is platform and configuration specific and should be characterized for best performance. Prior to DRAM training, BIOS should program this parameter to the largest defined value. After DRAM training, BIOS should use the guidelines below to configure the recommended platform generic timing values after DDR training is complete:

- $TrwtTO$ (in MEMCLKs) = $CEIL(MAX(ROD, CDD_{TrwtTO} / 2 - 0.5 + WrEarly) + LD + 3)$.
 - If 1 DIMM/ch, substitute $ROD = 0$ in the above equation.

The Critical Delay Difference (CDD) is the largest delay difference of the channel.

- Each delay difference is $D18F2x9C_x0000_00[2A:10]_dct[1:0]_{mp}[1:0][DqsRcvEnGrossDelay]$ minus $D18F2x9C_x0000_00[4A:30]_dct[1:0]_{mp}[1:0][WrDqsGrossDly]$.
- For CDD_{TrwtTO} , the subtraction terms are the delays of all chip selects within the same byte lane.

2.9.5.6.5 DRAM ODT Control

The following describes the general ODT behavior for various system configurations. In all cases, the processor ODT is off for writes and is on for reads:

- For one single or dual rank DIMM on a channel:
 - For writes, the ODT is on for the target rank.
 - For reads, the ODT is off for all ranks.
- For two single or dual rank DIMMs on a channel:
 - For writes, the ODT is on for the target rank of the target DIMM and also on for the first rank of the non-target DIMM.
 - For reads, the ODT is on for the first rank of the non-target DIMM.

BIOS configures the DIMM ODT behavior per chip select according to the DIMM population. In all cases, the processor ODT is off for writes and is on for reads. The ODT pin patterns for reads and writes are programmed using $D18F2x[234:230]_dct[1:0]$ and $D18F2x[23C:238]_dct[1:0]$, respectively.

BIOS also configures the ODT pin pattern used during write leveling by programming

D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][WrLvOdt]. BIOS programs WrLvOdt with the D18F2x[23C:238]_dct[1:0] value provided for writes to the rank targeted by training. See 2.9.5.9.1 [Write Levelization Training].

Table 31: DDR3 DIMM ODT Pattern

DIMM0 ¹	DIMM1 ¹	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
		D18F2x230	D18F2x234	D18F2x238	D18F2x23C
-	SR	0000_0000h	0000_0000h	0004_0000h	0000_0000h
-	DR	0000_0000h	0000_0000h	0804_0000h	0000_0000h
SR	-	0000_0000h	0000_0000h	0000_0001h	0000_0000h
DR	-	0000_0000h	0000_0000h	0000_0201h	0000_0000h
SR/DR	SR/DR	0101_0404h	0000_0000h	0905_0605h	0000_0000h

1. DIMM0: MEMCSx[1:0], MEMODTx[1,0].
 DIMM1: MEMCSx[3:2], MEMODTx[3,2].
 Population restrictions may apply. See 2.9.5.6.6 for details.

2.9.5.6.6 DRAM Address Timing and Output Driver Compensation Control

This section describes the settings required for programming the timing on the address pins, the CS/ODT pins, and the CKE pins. The following tables document the address timing and output driver settings on a per channel basis for DDR3 DIMM types. DIMM0 is the DIMM closest to the processor on that channel and DIMM1 is the DIMM farthest from the processor on that channel. DIMMs must be populated from farthest slot to closest slot to the processor on a per channel basis. Populations that are not shown in these tables are not supported. These tables document the optimal settings for motherboards which meet the relevant motherboard design guidelines.

- The presence of settings for higher speeds does not constitute support for those speeds. See 2.9 [DRAM Controllers (DCTs)] for an overview of the DIMM and memory bus speed support.

Table 32: BIOS recommendations for UDIMM address timings and output driver control for FM2 or FS1r2 Package

Condition					D18F2x94_dct[1:0]	RTT_Nom		RTT_Wr		D18F2x9C_x0000_0004_dct[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dct[1:0]_mp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1	SlowAccessMode						
1	667	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00112222h		
1	667	1.25, 1.35, 1.5	DR	-	0	120	Off	003B0000h	00112222h		
1	800	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00112222h		

Table 32: BIOS recommendations for UDIMM address timings and output driver control for FM2 or FS1r2 Package

Condition					D18F2x94_dcl[1:0]	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dcl[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dcl[1:0]_mp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1					
1	800	1.25, 1.35, 1.5	DR	-	0	120	Off	003B0000h	00112222h
1	1066	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	10112222h
1	1066	1.25, 1.35, 1.5	DR	-	0	120	Off	00380000h	10112222h
1	1333	1.25, 1.35, 1.5	SR	-	0	60	Off	00000000h	20112222h
1	1333	1.25, 1.35, 1.5	DR	-	0	60	Off	00360000h	20112222h
1	1600	1.35, 1.5	SR	-	0	60	Off	00000000h	30112222h
1	1600	1.35, 1.5	DR	-	1	40	Off	00000000h	30112222h
1	1866	1.5	SR	-	0	40	Off	00000000h	30112222h
1	1866	1.5	DR	-	1	40	Off	00000000h	30112222h
1	2133	1.5	SR	-	0	40	Off	00000000h	30112222h
1	2133	1.5	DR	-	1	40	Off	00000000h	30112222h
2	667	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00112222h
2	667	1.25, 1.35, 1.5	NP	DR	0	120	Off	003B0000h	00112222h
2	667	1.25, 1.35, 1.5	SR	SR	0	40	120	00390039h	10222322h
2	667	1.25, 1.35, 1.5	DR	DR	0	40	120	00390039h	10222322h
2	667	1.25, 1.35, 1.5	DR	SR	0	40	120	00390039h	10222322h
2	667	1.25, 1.35, 1.5	SR	DR	0	40	120	00390039h	10222322h
2	800	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00112222h
2	800	1.25, 1.35, 1.5	NP	DR	0	120	Off	003B0000h	00112222h
2	800	1.25, 1.35, 1.5	SR	SR	0	40	120	00390039h	20222322h
2	800	1.25, 1.35, 1.5	DR	DR	0	40	120	00390039h	20222322h
2	800	1.25, 1.35, 1.5	DR	SR	0	40	120	00390039h	20222322h
2	800	1.25, 1.35, 1.5	SR	DR	0	40	120	00390039h	20222322h
2	1066	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	10112222h
2	1066	1.25, 1.35, 1.5	NP	DR	0	120	Off	00380000h	10112222h
2	1066	1.25, 1.35, 1.5	SR	SR	0	40	120	00350037h	30222322h

Table 32: BIOS recommendations for UDIMM address timings and output driver control for FM2 or FS1r2 Package

Condition					D18F2x94_dcl[1:0]		D18F2x9C_x0000_0004_dcl[1:0]_mp[1:0]		D18F2x9C_x0000_0000_dcl[1:0]_mp[1:0]	
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1	SlowAccessMode	RTT		RTT_Nom	RTT_Wr	
						RTT_Nom	RTT_Wr			
2	1066	1.25, 1.35, 1.5	DR	DR	0	40	120	00350037h		30222322h
2	1066	1.25, 1.35, 1.5	DR	SR	0	40	120	00350037h		30222322h
2	1066	1.25, 1.35, 1.5	SR	DR	0	40	120	00350037h		30222322h
2	1333	1.25, 1.35, 1.5	NP	SR	0	60	Off	00000000h		20112222h
2	1333	1.25, 1.35, 1.5	NP	DR	0	60	Off	00360000h		20112222h
2	1333	1.25, 1.35, 1.5	SR	SR	1	30	120	00000035h		30222322h
2	1333	1.25, 1.35, 1.5	DR	DR	1	30	120	00000035h		30222322h
2	1333	1.25, 1.35, 1.5	DR	SR	1	30	120	00000035h		30222322h
2	1333	1.25, 1.35, 1.5	SR	DR	1	30	120	00000035h		30222322h
2	1600	1.35, 1.5	NP	SR	0	40	Off	00000000h		30112222h
2	1600	1.35, 1.5	NP	DR	1	40	Off	00000000h		30112222h
2	1600	1.35, 1.5	SR	SR	1	20	60	0000002Bh		30222322h
2	1600	1.35	DR	DR	1	20	60	0000002Bh		30222322h
2	1600	1.35	DR	SR	1	20	60	0000002Bh		30222322h
2	1600	1.35	SR	DR	1	20	60	0000002Bh		30222322h
2	1600	1.5	DR	DR	1	20	60	0000002Bh		30222322h
2	1600	1.5	DR	SR	1	20	60	0000002Bh		30222322h
2	1600	1.5	SR	DR	1	20	60	0000002Bh		30222322h
2	1866	1.5	NP	SR	0	40	Off	00000000h		30112222h
2	1866	1.5	NP	DR	1	40	Off	00000000h		30112222h
2	1866	1.5	SR	SR	1	20	60	00000031h		30222322h
2	1866	1.5	DR	DR	1	20	60	00000031h		30222322h
2	1866	1.5	DR	SR	1	20	60	00000031h		30222322h
2	1866	1.5	SR	DR	1	20	60	00000031h		30222322h
2	2133	1.5	NP	SR	0	40	Off	00000000h		30112222h
2	2133	1.5	NP	DR	1	40	Off	00000000h		30112222h

Table 33: BIOS recommendations for UDIMM address timings and output driver control for FP2 Package

Condition					D18F2x94_dcl[1:0]	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dcl[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dcl[1:0]_mp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1					
1	667	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00112222h
1	667	1.25, 1.35, 1.5	DR	-	0	120	Off	003B0000h	00112222h
1	800	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00112222h
1	800	1.25, 1.35, 1.5	DR	-	0	120	Off	003B0000h	00112222h
1	1066	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	10112222h
1	1066	1.25, 1.35, 1.5	DR	-	0	120	Off	00380000h	10112222h
1	1333	1.25, 1.35, 1.5	SR	-	0	60	Off	00000000h	20112222h
1	1333	1.25, 1.35, 1.5	DR	-	0	60	Off	00360000h	20112222h
2	667	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00112222h
2	667	1.25, 1.35, 1.5	NP	DR	0	120	Off	003B0000h	00112222h
2	667	1.25, 1.35, 1.5	SR	SR	0	40	120	00390039h	10222322h
2	667	1.25, 1.35, 1.5	DR	DR	0	40	120	00390039h	10222322h
2	667	1.25, 1.35, 1.5	DR	SR	0	40	120	00390039h	10222322h
2	667	1.25, 1.35, 1.5	SR	DR	0	40	120	00390039h	10222322h
2	800	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00112222h
2	800	1.25, 1.35, 1.5	NP	DR	0	120	Off	003B0000h	00112222h
2	800	1.25, 1.35, 1.5	SR	SR	0	40	120	00390039h	20222322h
2	800	1.25, 1.35, 1.5	DR	DR	0	40	120	00390039h	20222322h
2	800	1.25, 1.35, 1.5	DR	SR	0	40	120	00390039h	20222322h
2	800	1.25, 1.35, 1.5	SR	DR	0	40	120	00390039h	20222322h
2	1066	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	10112222h
2	1066	1.25, 1.35, 1.5	NP	DR	0	120	Off	00380000h	10112222h
2	1066	1.25, 1.35, 1.5	SR	SR	0	40	120	00350037h	30222322h
2	1066	1.25, 1.35, 1.5	DR	DR	0	40	120	00350037h	30222322h
2	1066	1.25, 1.35, 1.5	DR	SR	0	40	120	00350037h	30222322h

Table 33: BIOS recommendations for UDIMM address timings and output driver control for FP2 Package

Condition					D18F2x94_dctl[1:0]	SlowAccessMode	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dctl[1:0]_mnp[1:0]	D18F2x9C_x0000_0000_dctl[1:0]_mnp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1						
2	1066	1.25, 1.35, 1.5	SR	DR	0	40	120	00350037h	30222322h	
2	1333	1.25, 1.35, 1.5	NP	SR	0	60	Off	00000000h	20112222h	
2	1333	1.25, 1.35, 1.5	NP	DR	0	60	Off	00360000h	20112222h	
2	1333	1.25, 1.35, 1.5	SR	SR	1	30	120	00000035h	30222322h	
2	1333	1.35, 1.5	DR	DR	1	30	120	00000035h	30222322h	
2	1333	1.35, 1.5	DR	SR	1	30	120	00000035h	30222322h	
2	1333	1.35, 1.5	SR	DR	1	30	120	00000035h	30222322h	

Table 34: BIOS recommendations for SO-DIMM address timings and output driver control for FM2 or FS1r2 Package

Condition					D18F2x94_dctl[1:0]	SlowAccessMode	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dctl[1:0]_mnp[1:0]	D18F2x9C_x0000_0000_dctl[1:0]_mnp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1						
1	667	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00002222h	
1	667	1.25, 1.35, 1.5	DR	-	0	120	Off	00000000h	00002222h	
1	800	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00002222h	
1	800	1.25, 1.35, 1.5	DR	-	0	120	Off	00000000h	00002222h	
1	1066	1.25, 1.35, 1.5	SR	-	0	120	Off	003D3D3Dh	10002222h	

Table 34: BIOS recommendations for SO-DIMM address timings and output driver control for FM2 or FS1r2 Package

Condition					D18F2x94_dcl[1:0]	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dcl[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dcl[1:0]_mp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1					
1	1066	1.25, 1.35, 1.5	DR	-	0	120	Off	00000000h	10002222h
1	1333	1.25, 1.35, 1.5	SR	-	0	60	Off	003D3D3Dh	20002222h
1	1333	1.25, 1.35, 1.5	DR	-	0	60	Off	00003D3Dh	20002222h
1	1600	1.35, 1.5	SR	-	0	40	Off	003C3C3Ch	30112222h
1	1600	1.35, 1.5	DR	-	1	40	Off	00003C3Ch	30112222h
1	1866	1.5	SR	-	0	40	Off	003C3C3Ch	30112222h
1	1866	1.5	DR	-	1	40	Off	00003C3Ch	30112222h
2	667	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00002222h
2	667	1.25, 1.35, 1.5	NP	DR	0	120	Off	00000000h	00002222h
2	667	1.25, 1.35, 1.5	SR	SR	1	40	120	00000039h	10222323h
2	667	1.25, 1.35, 1.5	DR	DR	1	40	120	00000039h	10222323h
2	667	1.25, 1.35, 1.5	DR	SR	1	40	120	00000039h	10222323h
2	667	1.25, 1.35, 1.5	SR	DR	1	40	120	00000039h	10222323h
2	800	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00002222h
2	800	1.25, 1.35, 1.5	NP	DR	0	120	Off	00000000h	00002222h
2	800	1.25, 1.35, 1.5	SR	SR	1	40	120	00000039h	20222323h
2	800	1.25, 1.35, 1.5	DR	DR	1	40	120	00000039h	20222323h
2	800	1.25, 1.35, 1.5	DR	SR	1	40	120	00000039h	20222323h
2	800	1.25, 1.35, 1.5	SR	DR	1	40	120	00000039h	20222323h
2	1066	1.25, 1.35, 1.5	NP	SR	0	120	Off	003D3D3Dh	10002222h
2	1066	1.25, 1.35, 1.5	NP	DR	0	120	Off	00000000h	10002222h
2	1066	1.25, 1.35, 1.5	SR	SR	1	30	120	00000037h	30222323h
2	1066	1.25, 1.35, 1.5	DR	DR	1	30	120	00000037h	30222323h
2	1066	1.25, 1.35, 1.5	DR	SR	1	30	120	00000037h	30222323h
2	1066	1.25, 1.35, 1.5	SR	DR	1	30	120	00000037h	30222323h
2	1333	1.25, 1.35, 1.5	NP	SR	0	60	Off	003D3D3Dh	20002222h

Table 34: BIOS recommendations for SO-DIMM address timings and output driver control for FM2 or FS1r2 Package

Condition					D18F2x94_dcl[1:0]	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dcl[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dcl[1:0]_mp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1	SlowAccessMode				
2	1333	1.25, 1.35, 1.5	NP	DR	0	60	Off	00003D3Dh	20002222h
2	1333	1.25, 1.35, 1.5	SR	SR	1	30	120	00000035h	30222323h
2	1333	1.35, 1.5	DR	DR	1	30	120	00000035h	30222323h
2	1333	1.35, 1.5	DR	SR	1	30	120	00000035h	30222323h
2	1333	1.35, 1.5	SR	DR	1	30	120	00000035h	30222323h
2	1600	1.35, 1.5	NP	SR	0	40	Off	003C3C3Ch	30112222h
2	1600	1.35, 1.5	NP	DR	1	40	Off	00003C3Ch	30112222h
2	1600	1.5	SR	SR	1	20	60	00000033h	30222323h
2	1600	1.5	DR	DR	1	20	60	00000033h	30222323h
2	1600	1.5	DR	SR	1	20	60	00000033h	30222323h
2	1600	1.5	SR	DR	1	20	60	00000033h	30222323h
2	1866	1.5	NP	SR	0	40	Off	003C3C3Ch	30112222h
2	1866	1.5	NP	DR	1	40	Off	00003C3Ch	30112222h

Table 35: BIOS recommendations for SO-DIMM address timings and output driver control for FP2 Package

Condition					D18F2x94_dcl[1:0]	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dcl[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dcl[1:0]_mp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1					
1	667	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00002222h
1	667	1.25, 1.35, 1.5	DR	-	0	120	Off	00000000h	00002222h
1	800	1.25, 1.35, 1.5	SR	-	0	120	Off	00000000h	00002222h
1	800	1.25, 1.35, 1.5	DR	-	0	120	Off	00000000h	00002222h
1	1066	1.25, 1.35, 1.5	SR	-	0	120	Off	003D3D3Dh	10002222h
1	1066	1.25, 1.35, 1.5	DR	-	0	120	Off	00000000h	10002222h
1	1333	1.25, 1.35, 1.5	SR	-	0	60	Off	003D3D3Dh	20002222h
1	1333	1.25, 1.35, 1.5	DR	-	0	60	Off	00003D3Dh	20002222h
2	667	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00002222h
2	667	1.25, 1.35, 1.5	NP	DR	0	120	Off	00000000h	00002222h
2	667	1.25, 1.35, 1.5	SR	SR	1	40	120	00000039h	10222323h
2	667	1.25, 1.35, 1.5	DR	DR	1	40	120	00000039h	10222323h
2	667	1.25, 1.35, 1.5	DR	SR	1	40	120	00000039h	10222323h
2	667	1.25, 1.35, 1.5	SR	DR	1	40	120	00000039h	10222323h
2	800	1.25, 1.35, 1.5	NP	SR	0	120	Off	00000000h	00002222h
2	800	1.25, 1.35, 1.5	NP	DR	0	120	Off	00000000h	00002222h
2	800	1.25, 1.35, 1.5	SR	SR	1	40	120	00000039h	20222323h
2	800	1.25, 1.35, 1.5	DR	DR	1	40	120	00000039h	20222323h
2	800	1.25, 1.35, 1.5	DR	SR	1	40	120	00000039h	20222323h
2	800	1.25, 1.35, 1.5	SR	DR	1	40	120	00000039h	20222323h
2	1066	1.25, 1.35, 1.5	NP	SR	0	120	Off	003D3D3Dh	10002222h
2	1066	1.25, 1.35, 1.5	NP	DR	0	120	Off	00000000h	10002222h
2	1066	1.25, 1.35, 1.5	SR	SR	1	30	120	00000037h	30222323h
2	1066	1.25, 1.35, 1.5	DR	DR	1	30	120	00000037h	30222323h
2	1066	1.25, 1.35, 1.5	DR	SR	1	30	120	00000037h	30222323h

Table 35: BIOS recommendations for SO-DIMM address timings and output driver control for FP2 Package

Condition					D18F2x94_dct[1:0]	RTT_Nom		RTT_Wr		D18F2x9C_x0000_0004_dct[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dct[1:0]_mp[1:0]
NumDimmSlots	DdrRate	VDDIO	DIMM0	DIMM1	SlowAccessMode						
2	1066	1.25, 1.35, 1.5	SR	DR	1	30	120	00000037h			30222323h
2	1333	1.25, 1.35, 1.5	NP	SR	0	60	Off	003D3D3Dh			20002222h
2	1333	1.25, 1.35, 1.5	NP	DR	0	60	Off	00003D3Dh			20002222h
2	1333	1.25, 1.35, 1.5	SR	SR	1	30	120	00000035h			30222323h
2	1333	1.35, 1.5	DR	DR	1	30	120	00000035h			30222323h
2	1333	1.35, 1.5	DR	SR	1	30	120	00000035h			30222323h
2	1333	1.35, 1.5	SR	DR	1	30	120	00000035h			30222323h

Table 36: BIOS recommendations for DRAM soldered down address timings and output driver control for FP2 Package

Condition			D18F2x94_dct[1:0]	RTT_Nom		RTT_Wr		D18F2x9C_x0000_0004_dct[1:0]_mp[1:0]	D18F2x9C_x0000_0000_dct[1:0]_mp[1:0]
DdrRate	VDDIO	NumRanks	SlowAccessMode						
667	1.25, 1.35, 1.5	1	0	120	Off	00000000h			00002222h
667	1.25, 1.35, 1.5	2	0	120	Off	00000000h			00002222h
667	1.25, 1.35, 1.5	3	1	40	120	00000035h			30222323h
667	1.25, 1.35, 1.5	4	1	40	120	00000035h			30222323h
800	1.25, 1.35, 1.5	1	0	120	Off	00000000h			00002222h

Table 36: BIOS recommendations for DRAM soldered down address timings and output driver control for FP2 Package

Condition			D18F2x94_dct[1:0]	RTT_Nom	RTT_Wr	D18F2x9C_x0000_0004_dct[1:0]_mmp[1:0]	D18F2x9C_x0000_0000_dct[1:0]_mmp[1:0]
DdrRate	VDDIO	NumRanks	SlowAccessMode				
800	1.25, 1.35, 1.5	2	0	120	Off	00000000h	00002222h
800	1.25, 1.35, 1.5	3	1	40	120	00000035h	30222323h
800	1.25, 1.35, 1.5	4	1	40	120	00000035h	30222323h
1066	1.25, 1.35, 1.5	1	0	120	Off	003D3D3Dh	10002222h
1066	1.25, 1.35, 1.5	2	0	120	Off	00000000h	10002222h
1066	1.25, 1.35, 1.5	3	1	30	120	00000035h	30222323h
1066	1.25, 1.35, 1.5	4	1	30	120	00000035h	30222323h
1333	1.25, 1.35, 1.5	1	0	60	Off	003D3D3Dh	20002222h
1333	1.25, 1.35, 1.5	2	0	60	Off	00003D3Dh	20002222h
1333	1.25, 1.35, 1.5	3	1	30	120	00000035h	30222323h
1333	1.25, 1.35, 1.5	4	1	30	120	00000035h	30222323h

2.9.5.7 DCT Training Specific Configuration

The DCT requires certain features be disabled during DRAM device initialization and training. BIOS should program the registers in [Table 37](#) before DRAM device initialization and training. For normal operation, BIOS programs the recommended values if provided in [Table 37](#). BIOS must quiesce all other forms of DRAM traffic on the channel being trained. See [2.9.5 \[DCT/DRAM Initialization and Resume\]](#).

Table 37: DCT Training Specific Register Values

Register	Training	Normal Operation
D18F2x78_dct[1:0][AddrCmdTriEn]	0	1
D18F2x8C_dct[1:0][DisAutoRefresh]	1	0
D18F2x90_dct[1:0][ForceAutoPchg]	0	0
D18F2x90_dct[1:0][DynPageCloseEn]	0	0
D18F2x94_dct[1:0][BankSwizzleMode]	0	1
D18F2x94_dct[1:0][DcqBypassMax]	0	1Fh
D18F2x94_dct[1:0][PowerDownEn]	0	1
D18F2x94_dct[1:0][ZqcsInterval]	00b	10b

Table 37: DCT Training Specific Register Values

Register	Training	Normal Operation
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][RxMaxDurDllNoLock]	000b	See 2.9.5.11
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][TxMaxDurDllNoLock]	000b	See 2.9.5.11
D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0][EnRxPadStandby]	0	See 2.9.5.11
D18F2xA4[CmdThrottleMode]	000b	xxx ¹
D18F2xA4[ODTSEn]	000b	x ¹
D18F2xA4[BwCapEn]	0	x ¹
D18F2xA8_dct[1:0][BankSwap]	0	1
D18F2x110[DctSelIntLvEn]	0	x ¹
1. Programmed specific to the current platform or memory configuration.		

2.9.5.8 DRAM Device and Controller Initialization

BIOS initializes the DRAM devices and the controller using a software controlled sequence. See 2.9.5.8.1 [Software DDR3 Device Initialization].

BIOS must observe additional requirements for changing the PLL frequency when setting D18F2x7C_dct[1:0][EnDramInit]. See 2.9.5.4.2 [DRAM Channel Frequency Change].

DRAM initialization is complete after the value of D18F2x7C_dct[1:0][EnDramInit] is written by BIOS from 1 to 0 in the software-controlled sequence.

2.9.5.8.1 Software DDR3 Device Initialization

BIOS should apply the following procedure to each DCT to initialize the DDR3 DIMMs on the channel. This procedure should be run only when booting from an unpowered state (ACPI S4, S5 or G3; not S3, suspend to RAM).

1. Configure the DCT registers, including MemClkFreq and MemClkFreqVal.
2. Program D18F2x7C_dct[1:0][EnDramInit] = 1.
3. Wait 200 us.
4. Program D18F2x7C_dct[1:0][DeassertMemRstX] = 1.
5. Wait 500 us.
6. Program D18F2x7C_dct[1:0][AssertCke] = 1.
7. Wait 360 ns.
8. Send MRS(2).
9. Send MRS(3). Ordinarily at this time, MrsAddress[2:0] = 000b.
10. Send MRS(1) with MrsAddress[7] = 0.
11. Send MRS(0) with MrsAddress[8] = 1.
12. Send two ZQCL commands.
 - BIOS instructs the DCT to send a ZQCL command by programming D18F2x7C_dct[1:0] as follows:
 - Program MrsAddress[10] = 1.
 - Program SendZQCmd = 1.
 - Wait for SendZQCmd = 0.
 - Wait 512 MEMCLKs.

13. Program `D18F2x7C_dct[1:0][EnDramInit] = 0`.

14. Program `D18F2x2E8_dct[1:0]_mp[1:0]`, `D18F2x2EC_dct[1:0]_mp[1:0]` with a copy of the mode register data sent in the steps above, except with `D18F2x2E8_dct[1:0]_mp[1:0][8] = 0` (do not reset the DLL with a memory P-state change).

2.9.5.8.1.1 DDR3 MR Initialization

BIOS instructs the DCT to send MRS commands by programming `D18F2x7C_dct[1:0]` as follows:

1. Program MrsBank and MrsAddress as specified by Table 38, Table 39, Table 40, and Table 41:

- MrsBank[2:0] = BA2:BA0.
- MrsAddress[15:0] = A15:A0.
- See also `D18F2x[5C:40]_dct[1:0][OnDimmMirror]`.

2. Program MrsChipSel as appropriate.

3. Program SendMrsCmd = 1.

4. Wait for SendMrsCmd = 0.

Table 38. DDR3 MR0

Address Field	Field	Value																		
BA2:BA0	MR Select	000b																		
A15:A13	Reserved	000b																		
A12	PPD	See <code>D18F2x84_dct[1:0][PchgPDMoSel]</code> .																		
A11:A9	WR	<table border="0"> <tr> <td><u>Bits</u></td> <td><u>Twr</u> (See <code>D18F2x22C_dct[1:0]_mp[1:0][Twr]</code>)</td> </tr> <tr> <td>000b</td> <td>16 MEMCLKs</td> </tr> <tr> <td>001b</td> <td>5 MEMCLKs</td> </tr> <tr> <td>010b</td> <td>6 MEMCLKs</td> </tr> <tr> <td>011b</td> <td>7 MEMCLKs</td> </tr> <tr> <td>100b</td> <td>8 MEMCLKs</td> </tr> <tr> <td>101b</td> <td>10 MEMCLKs</td> </tr> <tr> <td>110b</td> <td>12 MEMCLKs</td> </tr> <tr> <td>111b</td> <td>14 MEMCLKs</td> </tr> </table>	<u>Bits</u>	<u>Twr</u> (See <code>D18F2x22C_dct[1:0]_mp[1:0][Twr]</code>)	000b	16 MEMCLKs	001b	5 MEMCLKs	010b	6 MEMCLKs	011b	7 MEMCLKs	100b	8 MEMCLKs	101b	10 MEMCLKs	110b	12 MEMCLKs	111b	14 MEMCLKs
<u>Bits</u>	<u>Twr</u> (See <code>D18F2x22C_dct[1:0]_mp[1:0][Twr]</code>)																			
000b	16 MEMCLKs																			
001b	5 MEMCLKs																			
010b	6 MEMCLKs																			
011b	7 MEMCLKs																			
100b	8 MEMCLKs																			
101b	10 MEMCLKs																			
110b	12 MEMCLKs																			
111b	14 MEMCLKs																			
A8	DLL	Controlled as required by the initialization sequence																		
A7	TM	0																		

Table 38. DDR3 MR0

Address Field	Field	Value	
A6:A4,A2	CAS Latency	Bits	<u>D18F2x200[Tcl]</u>
		0000b	Reserved
		0001b	12 MEMCLKs
		0010b	5 MEMCLKs
		0011b	13 MEMCLKs
		0100b	6 MEMCLKs
		0101b	14 MEMCLKs
		0110b	7 MEMCLKs
		0111b	15 MEMCLKs
		1000b	8 MEMCLKs
		1001b	16 MEMCLKs
		1010b	9 MEMCLKs
		1011b	Reserved
		1100b	10 MEMCLKs
		1101b	Reserved
		1110b	11 MEMCLKs
1111b	Reserved		
A3	RBT	1	
A1:A0	BL	<u>D18F2x84_dct[1:0]</u> [BurstCtrl]	

Table 39. DDR3 MR1

Address Field	Field	Value
BA2:BA0	MR Select	001b
A15:A13	Reserved	000b
A12	Qoff	0b
A11	TDQS	<u>D18F2x94_dct[1:0]</u> [RDqsEn]
A10	Reserved	0b
A8	Reserved	0b
A7	Level	Controlled as required by the initialization sequence
A4:A3	AL	00b
A9, A6, A2	Rtt_Nom	See 2.9.5.6.5 [DRAM ODT Control]
A5, A1	DIC	01b
A0	DLL	0

Table 40. DDR3 MR2

Address Field	Field	Value
BA2:BA0	MR Select	010b
A15:A11	Reserved	0_0000b
A10:A9	Rtt_Wr	See 2.9.5.6.5 [DRAM ODT Control]
A8	Reserved	0b

Table 40. DDR3 MR2

Address Field	Field	Value
A7	SRT	See ASR
A6	ASR	See DIMM SPD Byte 31: SDRAM Thermal and Refresh Options
A5:A3	CWL	D18F2x20C_dct[1:0]_mp[1:0][Tcwl]-5
A2:A0	PASR	000b

Table 41. DDR3 MR3

Address Field	Field	Value
BA2:BA0	MR Select	011b
A15:A3	Reserved	000_0000_0000b
A2	MPR	0b
A1:A0	MPR Loc	00b

2.9.5.9 DRAM Training

This section describes detailed methods used to train the processor DDR interface to DRAM for optimal functionality and performance. DRAM training is performed by BIOS after initializing the DRAM controller. See [2.9.5.8 \[DRAM Device and Controller Initialization\]](#).

Some of the DRAM training steps described in this section require two passes if the target MEMCLK frequency is greater than 333 MHz. For optimal software performance, software may defer the second pass (at target MEMCLK frequency) for each training step until after the first pass (at lowest supported frequency) of all other training steps are complete. See [D18F2x94_dct\[1:0\]\[MemClkFreq\]](#).

Product specific training requirements are as follows:

- Program [D18F2x9C_x0000_0008_dct\[1:0\]_mp\[1:0\]\[TrNibbleSel\]=0](#).
- Program [D18F2xA8_dct\[1:0\]\[PerRankTimingEn\]=1](#).

See [2.9.5.7 \[DCT Training Specific Configuration\]](#) for additional training requirements.

In the following subsections, lane is used to describe an 8-bit wide data group, each with its own timing control.

2.9.5.9.1 Write Levelization Training

Write levelization involves using the phy to detect the edge of DQS with respect to the memory clock on the DIMM for write accesses to each lane.

Training is accomplished on a per channel, per rank basis. If the target frequency is greater than 333 MHz then BIOS performs two passes; otherwise, only one pass is required. See [2.9.5.4.2 \[DRAM Channel Frequency Change\]](#).

- Pass 1: Configure the memory subsystem for 333 MHz MEMCLK frequency.
- Pass 2: Configure the memory subsystem for the target MEMCLK frequency.

The following describes the steps used for each pass of write levelization training for each channel:

For each rank:

1. Prepare the DIMMs for write levelization using DDR3-defined MR commands.
 - A. Configure the output driver and on-die termination of the target DIMM as follows:
 - For the target rank of the target DIMM, enable write leveling mode and enable the output driver.
 - For all other ranks of the target DIMM, enable write leveling mode and disable the output driver.
 - For two or more DIMMs per channel, program Rtt_Nom of the target rank to the corresponding specified Rtt_Wr termination. Otherwise, configure Rtt_Nom of the target DIMM as normal. See [2.9.5.6.6 \[DRAM Address Timing and Output Driver Compensation Control\]](#).
 - B. Configure Rtt_Nom on the non-target DIMMs as normal. See [2.9.5.6.6 \[DRAM Address Timing and Output Driver Compensation Control\]](#).
2. Wait 40 MEMCLKs.
3. Configure the phy for write levelization training:
 - A. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][WrtLvTrEn]=0`.
 - B. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][TrChipSel]` to specify the target rank to be trained.
 - C. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][WrLvOdt]` to the proper ODT settings for the current memory subsystem configuration. See [2.9.5.6.5 \[DRAM ODT Control\]](#).
 - D. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][WrLvOdtEn]=1`.
 - E. MFENCE.
 - F. Wait 10 MEMCLKs to allow for ODT signal settling.
 - G. For each lane program an initial value to registers `D18F2x9C_x0000_00[51:50]_dct[1:0]` to set the gross and fine delay. See [2.9.5.9.1.1 \[Write Leveling Seed Value\]](#).
4. Perform write leveling of the devices on the DIMM:
 - A. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][WrtLvTrEn]=1`.
 - B. MFENCE.
 - C. Wait 200 MEMCLKs.
 - D. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][WrtLvTrEn]=0`.
 - E. Read from registers `D18F2x9C_x0000_00[51:50]_dct[1:0]` to get the gross and fine delay settings for the target DIMM and save these values.
5. Disable write levelization training so that the phy stops driving write levelization ODT.
 - A. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][WrLvOdtEn]=0`.
 - B. MFENCE.
 - C. Wait 10 MEMCLKs to allow for ODT signal settling.
6. Program the target DIMM back to normal operation by configuring the following:
 - A. Configure all ranks of the target DIMM for normal operation.
 - B. Enable the output drivers of all ranks of the target DIMM.
 - C. For a two or more DIMM system, program the Rtt_Nom value for the target DIMM to the normal operating termination.

For each rank:

- BIOS calculates and programs the final saved gross and fine delay values for each lane into `D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0]` [[DRAM DQS Write Timing](#)].
 - $WrDqsFineDly = PhRecFineDlyByte$.
 - $GrossDly = SeedGross + PhRecGrossDlyByte - SeedPreGross$.
 - The Critical Gross Delay (CGD) is the minimum GrossDly of all byte lanes and all DIMMs.
 - If $(CGD < 0)$ Then
 - $D18F2x20C_dct[1:0]_mp[1:0][WrDqDqsEarly] = ABS(CGD)$
 - $WrDqsGrossDly = GrossDly + WrDqDqsEarly$
 - Else

- $D18F2x20C_dct[1:0]_mp[1:0][WrDqDqsEarly] = 0$.
- $WrDqsGrossDly = GrossDly$.

2.9.5.9.1.1 Write Leveling Seed Value

The seed value for pass 1 of write leveling is design and platform specific. The platform vendor may need to characterize and adjust this value for proper write levelization training. The seed delay value must fall within +/- 1/2 MEMCLK, including silicon PVT margin and jitter of the measured clock delay.

1. Calculate the total seed based on the following:
 - Pass 1: $SeedTotal =$ configuration specific seed value found in [Table 42](#).
 - Pass N:
 - $SeedTotalPreScaling =$ the total delay values obtained from the previous (N-1) pass of write levelization training.
 - $Write\ Leveling\ Total\ Delay = D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0][WrDqsGrossDly, WrDqs-FineDly] - (0x20 * D18F2x20C_dct[1:0]_mp[1:0][WrDqDqsEarly])$
 - $SeedTotal = SeedTotalPreScaling * (target\ frequency) / (frequency\ from\ previous\ pass)$.
2. $SeedGross = SeedTotal\ DIV\ 32$.
3. $SeedFine = SeedTotal\ MOD\ 32$.
4. If (SeedGross is odd)
 - then $SeedPreGross = 1$
 - else $SeedPreGross = 2$.
5. Program $D18F2x9C_x0000_00[51:50]_dct[1:0][PhRecFineDlyByte] = SeedFine$.
6. Program $D18F2x9C_x0000_00[51:50]_dct[1:0][PhRecGrossDlyByte] = SeedPreGross$.

Table 42. DDR3 Write Leveling Seed Values

DIMM Type	Seed Value ¹
SO-DIMM	0Eh
Unbuffered	15h
1. DDR3-667 (333 MHz).	

2.9.5.9.2 DQS Receiver Enable Training

Receiver enable delay training is used to dynamically determine the optimal delay value for [D18F2x9C_x0000_00\[2A:10\]_dct\[1:0\]_mp\[1:0\]](#) [[DRAM DQS Receiver Enable Timing](#)]. The optimal DQS receiver enable delay value is platform and load specific, and occurs in the middle of a received read preamble. The timing of the preamble includes the inbound DQS propagation delay, which is unknown by BIOS. The training for delay values involves:

1. Configuring the phy for an initial expected phase value (seed).
2. Generating a stream of read DQS edges from the DRAM by issuing multiple read commands.
3. The phy determining the phase between the received DQS edges and a reference clock.
4. Calculating a final delay value for enabling receivers during normal read operations using the phase determined by the phy.

BIOS should program $D18F2x210_dct[1:0]_nbp[3:0][MaxRdLatency]$ to 55h.

Training is accomplished on a per channel, per rank basis. If the target frequency is greater than 333 MHz then BIOS performs two passes; otherwise only one pass is required. See [2.9.5.4.2 \[DRAM Channel Frequency](#)

Change].

- Pass 1: Configure the memory subsystem for 333 MHz MEMCLK frequency.
- Pass 2: Configure the memory subsystem for the target MEMCLK frequency.

The following describes the steps used for each pass of receiver enable training for each channel:

For each rank:

1. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][TrChipSel]` to specify the target rank to be trained.
2. For each lane program an initial value to registers `D18F2x9C_x0000_00[51:50]_dct[1:0]` to set the gross and fine delay as specified in 2.9.5.9.2.1 [DQS Receiver Enable Training Seed Value].
3. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][DqsRcvTrEn]=1`.
4. Issue 192 read requests to the target rank. See 2.9.5.9.6 [Continuous Pattern Generation].
5. Program `D18F2x9C_x0000_0008_dct[1:0]_mp[1:0][DqsRcvTrEn]=0`.
6. Read `D18F2x9C_x0000_00[51:50]_dct[1:0][PhRecGrossDlyByte, PhRecFineDlyByte]` to get the gross and fine delay values for each lane.
7. For each lane, calculate and program the corresponding receiver enable delay values for `D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnGrossDelay, DqsRcvEnFineDelay]`. Save the result for use later.
 - $DqsRcvEnFineDelay = PhRecFineDlyByte$
 - $DqsRcvEnGrossDelay = SeedGross + PhRecGrossDlyByte - SeedPreGross + 1$.

2.9.5.9.2.1 DQS Receiver Enable Training Seed Value

The seed value for pass 1 of receiver enable delay training is design and platform specific and should be determined by characterization for best performance. The seed value represents the total delay from a reference point to the first expected rise edge of DQS on a read CAS measured at the processor pins, in 1 UI/32 increments. The reference point is defined as the clock in which CAS is asserted + CL - 1. This value is expected to be larger than 2 UI in the steps below.

The following steps are taken to determine the seed values needed to program the DRAM Phase Recovery Control Registers:

For each pass and each lane:

1. Calculate the total seed based on the following:
 - Pass 1: $SeedTotal = \text{The seed value found in Table 43} + \text{the total delay values obtained from the first pass of write levelization training. See 2.9.5.9.1 [Write Levelization Training].}$
 - $Write\ Leveling\ Total\ Delay = D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0][WrDqsGrossDly, WrDqs-FineDly] - (0x20 * D18F2x20C_dct[1:0]_mp[1:0][WrDqDqsEarly])$
 - Pass N:
 - $RegisterDelay = 0$
 - $SeedTotalPreScaling = (\text{the total delay values in } D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0] \text{ from the previous (N-1) pass of training}) - RegisterDelay - 20h.$
 - $SeedTotal = RegisterDelay + FLOOR(SeedTotalPreScaling * (\text{target frequency}) / (\text{frequency from previous pass}))$.
2. $SeedGross = SeedTotal \text{ DIV } 32$.
3. $SeedFine = SeedTotal \text{ MOD } 32$.
4. If (SeedGross is odd)
 - then $SeedPreGross = 1$
 - else $SeedPreGross = 2$.
5. Program `D18F2x9C_x0000_00[51:50]_dct[1:0][PhRecFineDlyByte] = SeedFine`.

6. Program `D18F2x9C_x0000_00[51:50]_dct[1:0][PhRecGrossDlyByte]` = SeedPreGross.
7. Program `D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnGrossDelay]` = SeedGross.

Table 43. DDR3 DQS Receiver Enable Training Seed Values

DIMM Type	Seed Value ¹
SO-DIMM	32h
Unbuffered	32h
1. DDR3-667 (333 MHz)	

2.9.5.9.3 DQS Receiver Enable Cycle Training

Receiver enable delay cycle training is used to train the gross delay settings of `D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0]` to the middle of the received read preamble using the phy phase results.

For each rank and lane:

1. Program `D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0][BlockRxDqsLock]` = 1.
2. $RxEnOrig = D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnGrossDelay, DqsRcvEnFineDelay]$ result from 2.9.5.9.2 [DQS Receiver Enable Training].
3. $RxEnOffset = MOD(RxEnOrig + 10h, 40h)$
4. For each DqsRcvEn value beginning from RxEnOffset incrementing by 1 MEMCLK:
 - A. Program `D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnGrossDelay, DqsRcvEnFineDelay]` with the current value.
 - B. Perform 2.9.5.9.4 [DQS Position Training].
 - Record the result for the current DqsRcvEn setting as a pass or fail depending if a data eye is found.
5. Process the array of results and determine a pass-to-fail transition.
 - A. $DqsRcvEnCycle =$ the total delay value of the pass result.
 - B. Program `D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnGrossDelay, DqsRcvEnFineDelay]` = $DqsRcvEnCycle - 10h$.
6. Program `D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0][BlockRxDqsLock]` = 0.

2.9.5.9.4 DQS Position Training

DQS position training is used to place the DQS strobe in the center of the read DQ data eye and to center the write DQ data eye across the write DQS strobe. Determining the correct DRAM DQS and DQ delay settings for both reads and writes is conducted by performing a two dimensional search of the delay settings found in `D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0]` [DRAM Read DQS Timing] and `D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0]` [DRAM Write Data Timing].

Training is accomplished on a per channel, per rank, and per lane basis.

For DQS position training, BIOS generates a training pattern using continuous read or write data streams. A 2k-bit-time training pattern is recommended for optimal results. To achieve this, BIOS programs `D18F2x260_dct[1:0][CmdCount]` = 256, `D18F2x250_dct[1:0][CmdTgt]`=01b, and `D18F2x25[8,4]_dct[1:0]` to access two different banks of the same CS. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Prior to DQS position training, BIOS must program `D18F2x210_dct[1:0]_nbp[3:0][MaxRdLatency]` based on

the current greatest value of `D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0]`. See 2.9.5.9.5 [Calculating MaxRdLatency].

The following describes the steps used for DQS position training for each channel:

- Program `D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0]` for all lanes.
 - `RxBypass3rd4thStg = 1`.
 - `Rx4thStgEn = 0`.

For each rank and lane:

1. Select a 64 byte aligned test address.
2. For each write data delay value in `D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0]` from Wr-DQS to Wr-DQS plus 1 UI, using the Wr-DQS delay value found in 2.9.5.9.1 [Write Levelization Training]:
 - Program the write data delay value for the current lane.
 - Write the DRAM training pattern to the test address.
3. For each read DQS delay value in `D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0]` from 0 to 1 UI:
 - Program the read DQS delay value for the current lane.
 - Read the DRAM training pattern from the test address.
 - Record the result for the current settings as a pass or fail depending if the pattern is read correctly.
 - Process the array of results and determine the longest string of consecutive passing read DQS delay values.
 - If the read DQS delay results for the current lane contain three or more consecutive passing delay values, then program `D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0]` with the average value of the smallest and largest delay values in the string of consecutive passing results.
 - If the average value of passing read DQS delay for the lane is negative, then adjust the input receiver DQ delay in `D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0]` for the lane as follows:
 - IF (`RxBypass3rd4thStg == 1`) program `RxBypass3rd4thStg=0` and repeat step 3 above for all ranks and lanes.
 - ELSEIF (`Rx4thStgEn == 0`) program `Rx4thStgEn=1` and repeat step 3 above for all ranks and lanes.
 - ELSE program the read DQS delay for the lane with a value of zero.
4. Process the array of results and determine the longest string of consecutive passing write data delay values for the read DQS delay value found in the step above.
 - If the write data delay results for the current lane contain three or more consecutive passing delay values, then program `D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0]` with the average value of the smallest and largest delay values in the string of consecutive passing results.

See Figure 3.

3. $P = P + (8 - \text{D18F2x210_dct}[1:0]_nbp[3:0][\text{RdPtrInit}])$
4. $P = P + 5$
5. $P = P + (2 * (\text{D18F2x200_dct}[1:0]_mp[1:0][\text{Tcl}] - 1 \text{ clocks}))$
6. $P = P + \text{CEIL}(\text{MAX}(\text{D18F2x9C_x0000_00}[2A:10]_dct[1:0]_mp[1:0][\text{DqsRcvEnGrossDelay}, \text{DqsRcvEnFineDelay}] + \text{D18F2x9C_x0000_0}[3:0]0[6:5]_dct[1:0]_mp[1:0][\text{RdDqsTime}] \text{ PCLKs})) + 1$
7. If $(\text{NclkFreq}/\text{MemClkFreq} < 2)$ then $P = P + 4.5$
Else $P = P + 2.5$
8. $T = T + 1050 \text{ ps}$
9. $N = (P/(\text{MemClkFreq} * 2) + T) * \text{NclkFreq}$
 - See [D18F5x1\[6C:60\]\[NbDid, NbFid\]](#) and [D18F2x88_dct\[1:0\]\[MemClkFreq\]](#).
10. $\text{D18F2x210_dct}[1:0]_nbp[3:0][\text{MaxRdLatency}] = \text{CEIL}(N) - 1$

2.9.5.9.5.1 MaxRdLatency Training

After DRAM DQS receiver enable training, BIOS optimizes [D18F2x210_dct\[1:0\]_nbp\[3:0\]\[MaxRdLatency\]](#) using the following algorithm. For MaxRdLatency training, BIOS generates a training pattern using continuous read or write data streams. See [2.9.5.9.6.1 \[DRAM Training Pattern Generation\]](#).

For each channel:

1. Calculate a starting MaxRdLatency delay value by executing the steps in [2.9.5.9.5](#), excluding steps 4, 7, 8, and 10.
2. Select 32 64-byte aligned test addresses associated with the rank that has the worst case $(\text{D18F2x9C_x0000_00}[2A:10]_dct[1:0]_mp[1:0][\text{DqsRcvEnGrossDelay}, \text{DqsRcvEnFineDelay}] + \text{D18F2x9C_x0000_0}[3:0]0[6:5]_dct[1:0]_mp[1:0][\text{RdDqsTime}])$ register setting.
3. Write the DIMM test addresses with the training pattern.
4. For each MaxRdLatency value incrementing from the value calculated in step 1:
 - A. Program [D18F2x210_dct\[1:0\]_nbp\[3:0\]\[MaxRdLatency\]](#) with the current value.
 - B. Read the DIMM test addresses.
 - C. Compare the values read against the pattern written.
 - If the pattern is read correctly, go to step 5.
 - D. Program [D18F2x9C_x0000_0050_dct\[1:0\]=00000000h](#)
5. Program [D18F2x210_dct\[1:0\]_nbp\[3:0\]\[MaxRdLatency\]](#) = $\text{CEIL}(\text{current value} + 1 \text{ NCLK} + 1.5 \text{ MEM-CLK})$.

2.9.5.9.6 Continuous Pattern Generation

DRAM training relies on the ability to generate a string of continuous reads or writes between the processor and DRAM, such that worst case electrical interactions can be created. This section describes how these continuous strings of accesses may be generated.

2.9.5.9.6.1 DRAM Training Pattern Generation

DRAM training pattern generation uses PRBS generators in the DCT to generate controlled read and write traffic streams. During write pattern generation, data values based off of the seeded PRBS are burst to the DRAM interface. Conversely for reads, data bursts from the DRAM interface are compared against expected data values based off of the seeded PRBS on a per nibble basis.

Two address modes are available for DRAM training pattern generation, as configured by [D18F2x250_dct\[1:0\]\[CmdTgt\]](#). For generating a stream of reads or writes to the same rank, address target A mode is used. To generate a stream of accesses to up to two different ranks, address target A and B mode is used.

An overview of the BIOS sequence to generate training patterns is as follows:

- Configure the DCT for pattern generation. See 2.9.5.7 [DCT Training Specific Configuration].
- Ensure DIMMs are configured to support 8-beat bursts (BL8 or dynamic burst length on the fly). See 2.9.5.8.1.1 [DDR3 MR Initialization].
- Wait for `D18F2x250_dct[1:0][CmdSendInProg] = 0`.
- Program `D18F2x250_dct[1:0][CmdTestEnable] = 1`.
- Send activate commands as appropriate. See 2.9.5.9.6.1.1 [Activate and Precharge Command Generation].
- Send read or write commands as desired. See 2.9.5.9.6.1.2 [Read and Write Command Generation].
- Send precharge commands as appropriate. See 2.9.5.9.6.1.1 [Activate and Precharge Command Generation].
- Program `D18F2x250_dct[1:0][CmdTestEnable] = 0`.

2.9.5.9.6.1.1 Activate and Precharge Command Generation

Prior to sending read or write commands, BIOS must send an activate command to a row in a particular bank of the DRAM devices for access. To send an activate command, BIOS performs the following steps:

- Program `D18F2x28C_dct[1:0]` to the desired address as follows:
 - `CmdChipSelect = CS[7:0]`.
 - `CmdBank = BA[2:0]`.
 - `CmdAddress = A[17:0]`.
- Program `D18F2x28C_dct[1:0][SendActCmd] = 1`.
- Wait until `D18F2x28C_dct[1:0][SendActCmd] = 0`.
- Wait 75 MEMCLKs.

After completing its accesses, BIOS must deactivate open rows in the DRAM devices. To send a precharge or precharge all command to deactivate open rows in a bank or in all banks, BIOS performs the following steps:

- Wait 25 MEMCLKs.
- Program `D18F2x28C_dct[1:0]` to the desired address as follows:
 - `CmdChipSelect = CS[7:0]`.
 - Precharge all command:
 - `CmdAddress[10] = 1`.
 - Precharge command:
 - `CmdAddress[10] = 0`.
 - `CmdBank = BA[2:0]`.
- Program `D18F2x28C_dct[1:0][SendPchgCmd] = 1`.
- Wait until `D18F2x28C_dct[1:0][SendPchgCmd] = 0`.
- Wait 25 MEMCLKs.

On an activate command, the LR-DIMM stores the address of the physical rank so it can direct subsequent CAS commands. Therefore, only one row per bank can be active for a logical rank of a LR-DIMM.

For LR-DIMMs, the behavior of precharge and precharge all commands issued to physical ranks associated with the logical rank selected by `D18F2x28C_dct[1:0][CmdChipSelect]` depends on the setting of F0RC14. In broadcast mode, the command is issued to each physical rank. In rank addressable mode, the command is issued to a specific physical rank BIOS specifies by programming `D18F2x28C_dct[1:0][CmdAddress[17:14]]`.

2.9.5.9.6.1.2 Read and Write Command Generation

BIOS performs the following steps for read pattern generation:

- Program `D18F2x278_dct[1:0]`, and `D18F2x274_dct[1:0]` with the data comparison masks for bit lanes of

interest.

- Program `D18F2x270_dct[1:0][DataPrbsSeed]` the seed for the desired PRBS.
- Program `D18F2x260_dct[1:0][CmdCount]` equal to the number of cache line commands.
- Program `D18F2x25[8,4]_dct[1:0]` to the initial address.
- Program `D18F2x250_dct[1:0]` as follows:
 - ResetAllErr and StopOnErr as desired. See 2.9.5.9.6.1.3 [PRBS Data Comparison].
 - CmdTgt corresponding to `D18F2x25[8,4]_dct[1:0]`.
 - CmdType = 000b.
 - SendCmd = 1.
- If `D18F2x260_dct[1:0][CmdCount] != 0` then
Wait for `D18F2x250_dct[1:0][TestStatus] = 1` and `D18F2x250_dct[1:0][CmdSendInProg] = 0`.
- Program `D18F2x250_dct[1:0][SendCmd] = 0`.
- Read `D18F2x264_dct[1:0]`, `D18F2x268_dct[1:0]`, and `D18F2x26C_dct[1:0]` if applicable.

BIOS performs the following steps for write pattern generation:

- Program `D18F2x270_dct[1:0][DataPrbsSeed]` the seed for the desired PRBS.
- Program `D18F2x260_dct[1:0][CmdCount]` equal to the number of cache line commands desired.
- Program `D18F2x25[8,4]_dct[1:0]` to the initial address.
- Program `D18F2x250_dct[1:0]` as follows:
 - CmdTgt corresponding to `D18F2x25[8,4]_dct[1:0]`.
 - CmdType = 001b.
 - SendCmd = 1.
- If `D18F2x260_dct[1:0][CmdCount] != 0` then
Wait for `D18F2x250_dct[1:0][TestStatus] = 1` and `D18F2x250_dct[1:0][CmdSendInProg] = 0`.
- Program `D18F2x250_dct[1:0][SendCmd] = 0`.

BIOS combines the two sets of steps listed above for alternating write and read pattern generation.

2.9.5.9.6.1.3 PRBS Data Comparison

The DCT compares the incoming read data against the expected PRBS sequence specified by `D18F2x270_dct[1:0][DataPrbsSeed]` during pattern generation. BIOS may choose to continue command generation and accumulate errors or stop command generation on the first error occurrence by programming `D18F2x250_dct[1:0][StopOnErr]`. Error information is reported via `D18F2x264_dct[1:0]`, `D18F2x268_dct[1:0]`, `D18F2x26C_dct[1:0]`, `D18F2x294_dct[1:0]`, and `D18F2x298_dct[1:0]` can be masked on per-bit basis by programming `D18F2x274_dct[1:0]` and `D18F2x278_dct[1:0]`. BIOS resets the error information by programming `D18F2x250_dct[1:0][ResetAllErr]=1`.

Error information is only valid in certain modes of `D18F2x250_dct[1:0][CmdType, CmdTgt]` and `D18F2x260_dct[1:0][CmdCount]` and when using 64 byte aligned addresses in `D18F2x25[8,4]_dct[1:0][TgtAddress]`. Some modes require a series of writes to setup a DRAM data pattern. See Table 44.

Table 44. Command Generation and Data Comparison

Commands	CmdType	CmdTgt	Maximum CmdCount
Read	000b	00b ¹	128
		01b ¹	256 ²
Write-Read	010b	00b	Infinite
		01b ³	256 ²

1. Requires setup writes to store a data pattern in DRAM. The write commands are generated using the same CmdTgt, CmdCount, and DataPrbsSeed settings.

2. D18F2x254[TgtAddress] != D18F2x258[TgtAddress].

3. Requires setup writes to store a data pattern in DRAM. The write commands are generated programming D18F2x254[TgtAddress] to the intended Target B, CmdTgt=00b, CmdCount to 1/2 of the intended command count, and the same DataPrbsSeed setting.

2.9.5.10 DRAM Channel Disable

The following steps are performed to disable an unused DRAM channel:

For the channel to be disabled:

1. Program D18F2x9C_x0000_000C_dct[1:0][CKETri] = 1111b.
2. Wait 24 MEMCLKs.
3. Program D18F2x94_dct[1:0][DisDramInterface] = 1.
4. If channel 0 is disabled, then program D18F2x9C_x0D0F_E018_dct[0][PhyPSMasterChannel] = 1.

2.9.5.11 DRAM Phy Power Savings

For power savings, BIOS should perform the following actions for each channel:

1. Program D18F2x88_dct[1:0][MemClkDis] to disable unused MEMCLK pins.
2. Program D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0][PwrDn] for unused MEMCLK pairs.
3. Program D18F2x9C_x0000_000C_dct[1:0][CKETri, ODTTri, ChipSelTri] to disable unused pins.
4. Program D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0][DIIIDisEarlyU] = 1.
5. then Program D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0][DIIIDisEarlyL] = 1.
6. D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0][RxDqsUDIIIPowerDown] = 1.
7. D18F2x9C_x0D0F_812F_dct[1:0][PARTri] = 1.
8. D18F2x9C_x0D0F_812F_dct[1:0][Add17Tri, Add16Tri] = {1b, 1b}.
9. IF (DimmsPopulated == 1) && ((D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CkeDrvStren]==010b) || (D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CkeDrvStren]==011b)) THEN
 program D18F2x9C_x0D0F_C000_dct[1:0]_mp[1:0][LowPowerDrvStrengthEn] = 1
 ELSE program D18F2x9C_x0D0F_C000_dct[1:0]_mp[1:0][LowPowerDrvStrengthEn] = 0 ENDIF.
10. Program D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0][EnRxPadStandby] = IF
 (D18F2x94_dct[1:0][MemClkFreq] <= 800 MHz) THEN 1 ELSE 0 ENDIF.
11. Program D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] as follows:
 - If (DDR rate <= 1600) TxMaxDurDIINoLock = RxMaxDurDIINoLock = 8h
 else TxMaxDurDIINoLock = RxMaxDurDIINoLock = 7h.
 - TxCPUpdPeriod = RxCPUpdPeriod = 011b.

- TxDLLWakeupTime = RxDLLWakeupTime = 11b.
12. Program `D18F2x248_dct[1:0]_mp[1:0]` and then `D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0]` as follows:
- For M1 context program `RxChMntClkEn=RxSsbMntClkEn=0`.
 - For M0 context program `RxChMntClkEn=RxSsbMntClkEn=1`.

2.9.6 Memory Interleaving Modes

The processor supports two different types of memory interleaving modes:

- Chip select: interleaves the physical address space over multiple DIMM ranks on a channel, as opposed to each DIMM owning single consecutive address spaces. See 2.9.6.1 [Chip Select Interleaving].
- Channel: interleaves the physical address space over multiple channels, as opposed to each channel owning single consecutive address spaces. See 2.9.6.2 [Channel Interleaving].

Any combination of these interleaving modes may be enabled concurrently.

2.9.6.1 Chip Select Interleaving

The chip select memory interleaving mode has the following requirements:

- The number of chip selects interleaved is a power of two.
- The chip selects are the same size and type.

A BIOS algorithm for programming `D18F2x[5C:40]_dct[1:0]` [DRAM CS Base Address] and `D18F2x[6C:60]_dct[1:0]` [DRAM CS Mask] in memory interleaving mode is as follows:

1. Program all DRAM CS Base Address and DRAM CS Mask registers using contiguous normalized address mapping.
2. For each enabled chip select, swap the corresponding `BaseAddr[38:27]` bits with the `BaseAddr[21:11]` bits as defined in Table 45.
3. For each enabled chip select, swap the corresponding `AddrMask[38:27]` bits with the `AddrMask[21:11]` bits as defined in Table 45.

Table 45. DDR3 Swapped Normalized Address Lines for CS Interleaving

DIMM Address Map ¹	Chip Select Size	(BankSwap, DctSelIntLvAddr) ²	Swapped Base Address and Address Mask bits	
			4 way CS interleaving	2 way CS interleaving
0001b	256-MB	0b,N/A	[29:28] and [17:16]	[28] and [16]
		1b,100b	[29:28] and [12:11]	[28] and [11]
		1b,101b	[29:28] and [13:12]	[28] and [12]
0010b	512-MB	0b,N/A	[30:29] and [17:16]	[29] and [16]
		1b,100b	[30:29] and [12:11]	[29] and [11]
		1b,101b	[30:29] and [13:12]	[29] and [12]
0101b	1-GB	0b,N/A	[31:30] and [17:16]	[30] and [16]
		1b,100b	[31:30] and [12:11]	[30] and [11]
		1b,101b	[31:30] and [13:12]	[30] and [12]

Table 45. DDR3 Swapped Normalized Address Lines for CS Interleaving

DIMM Address Map ¹	Chip Select Size	(BankSwap, DctSelIntLvAddr) ²	Swapped Base Address and Address Mask bits	
			4 way CS interleaving	2 way CS interleaving
0111b	2-GB	0b,N/A	[32:31] and [17:16]	[31] and [16]
		1b,100b	[32:31] and [12:11]	[31] and [11]
		1b,101b	[32:31] and [13:12]	[31] and [12]
1010b	4-GB	0b,N/A	[33:32] and [17:16]	[32] and [16]
		1b,100b	[33:32] and [12:11]	[32] and [11]
		1b,101b	[33:32] and [13:12]	[32] and [12]
1011b	8-GB	0b,N/A	[34:33] and [18:17]	[33] and [17]
		1b,100b	[34:33] and [12:11]	[33] and [11]
		1b,101b	[34:33] and [13:12]	[33] and [12]

1. See [D18F2x80_dct\[1:0\]](#) [DRAM Bank Address Mapping].

2. See [D18F2xA8_dct\[1:0\]](#)[BankSwap] and [D18F2x110](#)[DctSelIntLvAddr]. If Bank-Swap==1 && ([D18F2x110](#)[DctSelIntLvEn]==0 || DctSelIntLvAddr!=100b) then software swaps the bits shown in the rows for address bit 9 channel interleaving.

The following is an example of interleaving a 64-bit interface to DDR3 DRAM. The DRAM memory consists of two 512 MB dual rank DDR3 DIMMs.

- The register settings for contiguous memory mapping are:
 - [D18F2x80_dct\[1:0\]](#) = 0000_0011h. // CS0/1 = 256 MB; CS2/3 = 256 MB
 - [D18F2x40](#) = 0000_0001h. // 0 MB base
 - [D18F2x44](#) = 0010_0001h. // 256 MB base = 0 MB + 256 MB
 - [D18F2x48](#) = 0020_0001h. // 512 MB base = 256 MB + 256 MB
 - [D18F2x4C](#) = 0030_0001h. // 768 MB base = 512 MB + 256 MB
 - [D18F2x60](#) = 0008_FFE0h. // CS0/CS1 = 256 MB
 - [D18F2x64](#) = 0008_FFE0h. // CS2/CS3 = 256 MB
- The base address bits to be swapped are defined in [Table 45](#), 256MB chip select size, 4 way CS interleaving column. The BaseAddr[29:28] bits are specified by [D18F2x\[5C:40\]_dct\[1:0\]\[21:20\]](#). The BaseAddr[17:16] bits are specified by [D18F2x\[5C:40\]_dct\[1:0\]\[11:10\]](#).
 - [D18F2x40](#) = 0000_0001h.
 - [D18F2x44](#) = 0000_0401h.
 - [D18F2x48](#) = 0000_0801h.
 - [D18F2x4C](#) = 0000_0C01h.
- The AddrMask bits to be swapped are the same as the BaseAddr bits defined in the previous step. The AddrMask[29:28] bits are specified by [D18F2x\[6C:60\]_dct\[1:0\]\[21:20\]](#). The AddrMask[17:16] bits are specified by [D18F2x\[6C:60\]_dct\[1:0\]\[11:10\]](#).
 - [D18F2x60](#) = 0038_F3E0h.
 - [D18F2x64](#) = 0038_F3E0h.
- If BankSwap is enabled and DCT channel interleaving is enabled on system address bit 8, then the Base and AddrMask bits to be swapped are as follows:
 - [D18F2x40](#) = 0000_0001h.
 - [D18F2x44](#) = 0000_0021h.

- D18F2x48 = 0000_0041h.
 - D18F2x4C = 0000_0061h.
 - D18F2x60 = 0038_FF90h.
 - D18F2x64 = 0038_FF90h.
5. If BankSwap is enabled and DCT channel interleaving is enabled on system address bit 9, then the Base and AddrMask bits to be swapped are as follows:
- D18F2x40 = 0000_0001h.
 - D18F2x44 = 0000_0041h.
 - D18F2x48 = 0000_0081h.
 - D18F2x4C = 0000_00C1h.
 - D18F2x60 = 0038_FF20h.
 - D18F2x64 = 0038_FF20h.

2.9.6.2 Channel Interleaving

The channel memory interleaving mode requires that DIMMs are present on both channels. Channel interleaving is enabled by programming [D18F2x110\[DctSelIntLvEn\]](#) and [D18F2x110\[DctSelIntLvAddr\]](#) to specify how interleaving is performed between the DCTs. If the channels do not have the same amount of DRAM, [D18F2x110\[DctSelBaseAddr, DctSelHi, DctSelHiRngEn\]](#) are used to configure the interleaved region. See also [2.9.7 \[Memory Hoisting\]](#).

2.9.7 Memory Hoisting

Memory hoisting reclaims the otherwise inaccessible DRAM that would naturally reside in memory regions used by MMIO. When memory hoisting is configured by BIOS, DRAM physical addresses are repositioned above the 4 GB address level in the address map. In operation, the physical addresses are remapped in hardware to the normalized addresses used by a DCT.

The region of DRAM that is hoisted is defined to be from [D18F1xF0\[DramHoleBase\]](#) to the 4 GB level or from [FD_0000_0000h](#) to the 1 TB level. Hoisting is enabled by programming [D18F1xF0 \[DRAM Hole Address\]](#) and configuring the DCTs per the equations in this section.

[DramHoleSize](#) is defined in order to simplify the following equations in this section and is calculated as follows:

- Define the DRAM hole region as $\text{DramHoleSize}[31:24] = 100\text{h} - \text{D18F1xF0}[\text{DramHoleBase}[31:24]]$.

2.9.7.1 DramHoleOffset Programming

[D18F1xF0\[DramHoleOffset\]](#) is programmed based on the scenarios shown in [Figure 5](#):

- **Case 1:** If [D18F2x110\[DctSelHiRngEn\]](#) = 0 OR [D18F2x110\[DctSelHiRngEn\]](#) = 1 AND [DctSelBaseAddr](#) > [DramHoleBase](#) then:
 $\text{DramHoleOffset}[31:23] = \{\text{DramHoleSize}[31:24], 0\text{b}\} + \{\text{DramBaseAddr}[31:27], 0000\text{b}\};$
- **Case 2:** If [D18F2x110\[DctSelIntLvEn\]](#) = 0 AND [D18F2x110\[DctSelHiRngEn\]](#) = 1 AND [DctSelBaseAddr](#) < [DramHoleBase](#) then:
 $\text{DramHoleOffset}[31:23] = \{\text{DramHoleSize}[31:24], 0\text{b}\} + \{\text{DctSelBaseAddr}[31:27], 0000\text{b}\};$
- **Case 3:** If [D18F2x110\[DctSelIntLvEn\]](#) = 1 AND [D18F2x110\[DctSelHiRngEn\]](#) = 1 AND [DctSelBaseAddr](#) < [DramHoleBase](#) then
 $\text{DramHoleOffset}[31:23] = \{\text{DramHoleSize}[31:24], 0\text{b}\} + \{\text{DramBaseAddr}[31:27], 0000\text{b}\} + \{0\text{b}, (\text{DctSelBaseAddr}[31:27] - \text{DramBaseAddr}[31:27]), 000\text{b}\};$

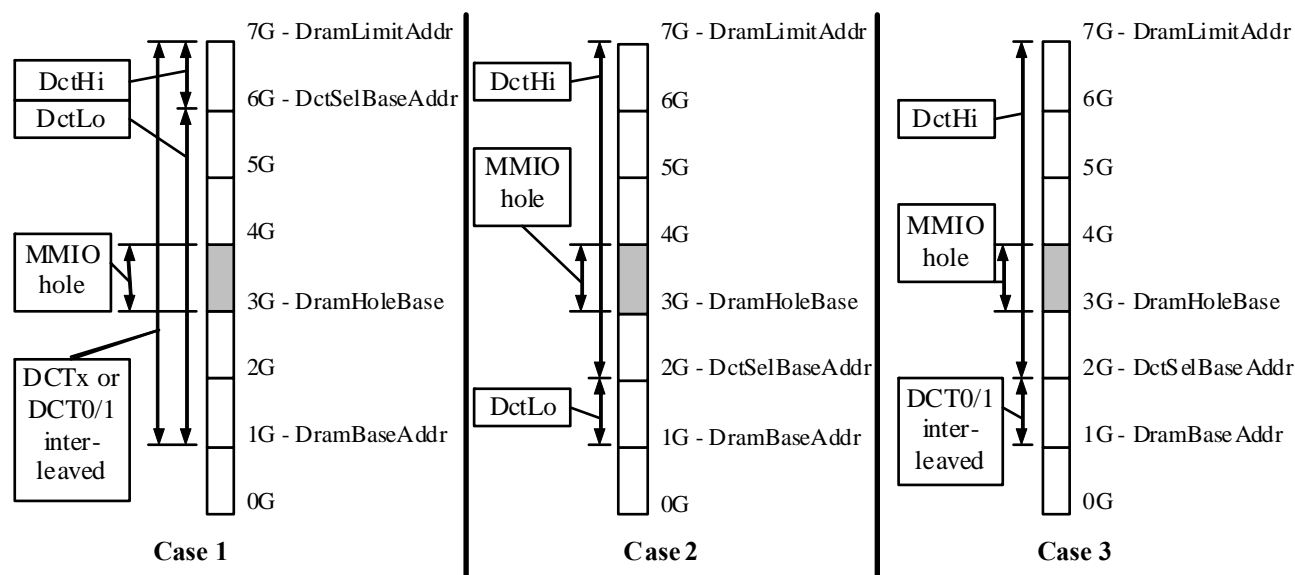


Figure 5: Example Cases for Programming DramHoleOffset

2.9.7.2 DctSelBaseOffset Programming

When $D18F2x110[DctSelHiRngEn] = 1$, $D18F2x114[DctSelBaseOffset]$ is programmed based on the scenarios shown in Figure 6:

- **Case 1:** If $D18F2x110[DctSelIntLvEn] = 0$ then:
 $DctSelBaseOffset[47:26] = \{DctSelBaseAddr[47:27], 0b\};$
- **Case 2:** If $D18F2x110[DctSelIntLvEn] = 1$ AND ($D18F1xF0[DramHoleValid] = 1$ AND $DctSelBaseAddr < DramHoleBase$ OR $D18F1xF0[DramHoleValid] = 0$) then:
 $DctSelBaseOffset[47:26] = \{DramBaseAddr[47:27], 0b\}$
 $+ \{0b, (DctSelBaseAddr[47:27] - DramBaseAddr[47:27])\};$
- **Case 3:**
 - If $D18F2x110[DctSelIntLvEn] = 1$ AND $D18F1xF0[DramHoleValid] = 1$ AND $DctSelBaseAddr > DramHoleBase$ then:
 $DctSelBaseOffset[47:26] = \{DramBaseAddr[47:27], 0b\}$
 $+ \{0000h, DramHoleSize[31:26]\}$
 $+ \{0b, (DctSelBaseAddr[47:27]$
 $- \{0000h, (DramBaseAddr[31:27] + DramHoleSize[31:27])\}\};$

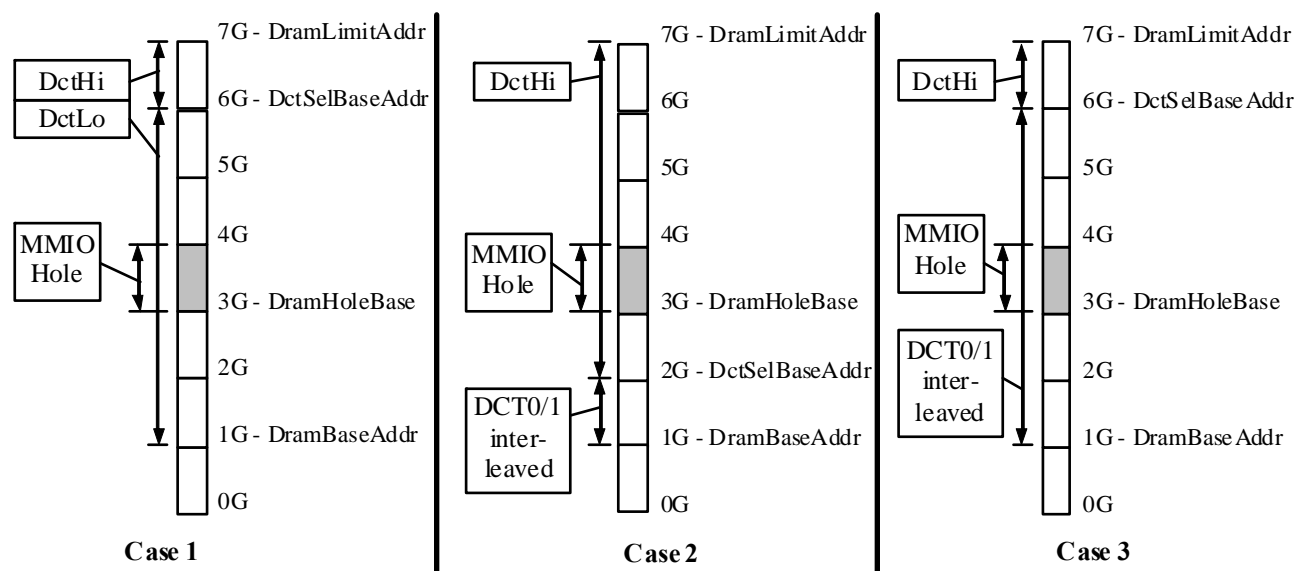


Figure 6: Example Cases for Programming DctSelBaseOffset

2.9.8 DRAM CC6/PC6 Storage

DRAM is used to hold the state information of cores entering the CC6 power management state. As part of the system setup if CC6 or PC6 is enabled, BIOS configures a special region of DRAM to hold the state information. In operation, hardware protects this region from general system accesses while allowing the cores access during C-state transitions.

The size of each special DRAM storage region is defined to be a fixed 16MB.

BIOS configures the storage region at the top of the DRAM range and adjusts `D18F1x[144:140,44:40][DramLimit]` downward accordingly.

See [Table 46](#).

After finalizing the system DRAM configuration, BIOS must set `D18F2x118[LockDramCfg] = 1` to enable the hardware protection.

Table 46. Example storage region configuration

Node	DRAM Populated	<code>D18F1x[144:140,44:40][DramBase, DramLimit]</code>	CC6 DRAM Range	<code>D18F4x128[CoreStateSaveDestNode]</code>	<code>D18F1x120[DramBaseAddr], D18F1x124[DramLimitAddr]</code>
0	256 MB	0 MB, 240 MB - 1	240 MB, 256 MB - 1	0	0 MB, 256 MB - 1

2.9.9 DRAM On DIMM Thermal Management and Power Capping

Each DCT can throttle commands based on the state of the channel `EVENT_L` pin or when `D18F2xA4[BwCapEn]=1`. The `EVENT_L` pin is used for thermal management while `D18F2xA4[BwCapEn]` limits memory power independent of the thermal management solution.

The EVENT_L pin for each channel must be wire OR'ed. If two DCTs are enabled then both throttle commands in lockstep using the amount specified in [D18F2xA4\[CmdThrottleMode\]](#) and [D18F2xA4\[BwCapCmdThrottleMode\]](#).

The recommended BIOS configuration for the EVENT_L pin is as follows:

- BIOS may enable command throttling on a DRAM controller if the platform supports the EVENT_L pin by programming [D18F2xA4\[ODTSEn\]](#) = 1.
 - The recommended usage is for this pin to be connected to one or more JEDEC defined on DIMM temperature sensors. The DIMM SPD ROM indicates on DIMM temperature sensor support.
 - BIOS configures the temperature sensor(s) to assert EVENT_L pin active low when the trip point is exceeded and deassert EVENT_L when the temperature drops below the trip point minus the sensor defined hysteresis.
 - BIOS programs [D18F2xA4\[CmdThrottleMode\]](#) with the throttling mode to employ when the trip point has been exceeded.
 - The hardware enforces a refresh rate of 3.9 us while EVENT_L is asserted.
- BIOS configures [D18F2x8C_dct\[1:0\]\[Tref\]](#) based on JEDEC defined temperature range options, as indicated by the DIMM SPD ROM. The two defined temperature ranges are normal (with a case temperature of 85 °C) and extended (with a case temperature of 95 °C).
 - If all DIMMs support the normal temperature range, or if normal and extended temperature range DIMMs are mixed, BIOS programs [D18F2x8C_dct\[1:0\]\[Tref\]](#) to 7.8 us and [D18F2xA4\[ODTSEn\]](#) = 1. BIOS configures the temperature sensor trip point for all DIMMs according to the 85 °C case temperature specification.
 - If all DIMMs support the extended temperature range, BIOS has two options:
 - a. Follow the recommendation for normal temperature range DIMMs.
 - b. Program [D18F2x8C_dct\[1:0\]\[Tref\]](#) = 3.9 us and configure the temperature sensor trip point for all DIMMs according to the 95 °C case temperature specification.
- At startup, the BIOS determines if the DRAMs are hot before enabling a DCT and delays for an amount of time to allow the devices to cool under the influence of the thermal solution. This is accomplished by checking the temperature status in the temperature sensor of each DIMM.
- The latched status of the EVENT_L pin for both DCTs can be read by system software in [D18F2xAC\[MemTempHot1, MemTempHot0\]](#).

The relationship between the DRAM case temperature, trip point, and EVENT_L pin sampling interval is outlined as follows:

- The trip point for each DIMM is ordinarily configured to the case temperature specification minus a guardband temperature for the DIMM.
- The temperature guardband is vendor defined and is used to account for sensor inaccuracy, EVENT_L pin sample interval, and platform thermal design.
- The sampling interval is vendor defined. It is expected to be approximately 1 second.

BIOS may enable bandwidth capping on a DRAM controller by setting [D18F2xA4\[BwCapEn\]](#) = 1 and programming [D18F2xA4\[BwCapCmdThrottleMode\]](#) with the throttling mode to employ. The DCT will employ the larger of the two throttling percentages as specified by [D18F2xA4\[BwCapCmdThrottleMode\]](#) and [D18F2xA4\[CmdThrottleMode\]](#) if the EVENT_L pin is asserted when both [D18F2xA4\[BwCapEn\]](#) = 1 and [D18F2xA4\[ODTSEn\]](#) = 1.

2.10 Thermal Functions

Thermal functions SB-TSI, HTC, PROCHOT_L and THERMTRIP are intended to maintain processor temperature in a valid range by:

- Providing a signal to external circuitry for system thermal management like fan control.
- Lowering power consumption by switching to lower-performance P-state.
- Sending processor to the THERMTRIP state to prevent it from damage.

The processor thermal-related circuitry includes (1) the temperature calculation circuit (TCC) for determining the temperature of the processor and (2) logic that uses the temperature from the TCC.

2.10.1 The Tctl Temperature Scale

Tctl is a processor temperature control value used for processor thermal management. Tctl is accessible through SB-TSI and `D18F3xA4[CurTmp]`. Tctl is a temperature on its own scale aligned to the processors cooling requirements. Therefore Tctl does not represent a temperature which could be measured on the die or the case of the processor. Instead, it specifies the processor temperature relative to the maximum operating temperature, Tctl,max. Tctl is defined as follows for all parts:

A: For $Tctl = Tctl_max$ to 255.875: the temperature of the part is $[Tctl - Tctl_max]$ over the maximum operating temperature. The processor may take corrective actions that affects performance, such as HTC, to support the return to Tctl range A.

B: For $Tctl = 0$ to $Tctl_max - 0.125$: the temperature of the part is $[Tctl_max - Tctl]$ under the maximum operating temperature.




Figure 7: Tctl scale

2.10.2 Temperature Slew Rate Control

The temperature slew rate controls in [D18F3xA4](#) are used to filter processor the processor temperature provided in [D18F3xA4\[CurTmp\]](#) and through SB-TSI. Separate controls are provided for increasing and decreasing temperatures. The latest measured temperature is referred to as Tctlm below.

If downward slew control is enabled ([D18F3xA4\[TmpSlewDnEn\]](#)), Tctl is not updated down unless Tctlm remains below Tctl for a time specified by [D18F3xA4\[PerStepTimeDn\]](#). If at any point before the timer expires Tctlm equals or exceeds Tctl, then the timer resets and Tctl is not updated. If the timer expires, then Tctl is reduced by 0.125. If downward slew control is disabled, then if Tctlm is less than Tctl, Tctl is immediately updated to Tctlm.

The upward slew control works similar to downward slew control except that if Tctlm exceeds Tctl by a value defined by [D18F3xA4\[TmpMaxDiffUp\]](#) then Tctl is immediately updated to Tctlm. Otherwise, Tctlm must remain above Tctl for time specified by [D18F3xA4\[PerStepTimeUp\]](#) before Tctl is incremented by 0.125.

2.10.3 Sideband Temperature Sensor Interface (SB-TSI)

SB-TSI is used by an external SMBus master to access the internal temperature sensor and to specify temperature thresholds. 100 kHz standard-mode and 400 kHz fast-mode are supported. 3.4 MHz high-speed mode is not supported.

2.10.4 Temperature-Driven Logic

The temperature calculated by the TCC is used by HTC, THERMTRIP, PROCHOT_L, and the serial interface, SB-TSI.

2.10.4.1 PROCHOT_L and Hardware Thermal Control (HTC)

The processor *HTC-active state* is characterized by (1) the assertion of PROCHOT_L, (2) reduced power consumption, and (3) reduced performance. While in the HTC-active state, the processor reduces power consumption by limiting all cores to a P-state (specified by [D18F3x64\[HtcPstateLimit\]](#)). See [2.5.3 \[CPU Power Management\]](#). While in the HTC-active state, software should not change the following: All [D18F3x64](#) fields (except for HtcActSts and HtcEn). Any change to the previous list of fields when in the HTC-active state can result in undefined behavior. HTC status and control is provided through [D18F3x64](#).

The PROCHOT_L pin acts as both an input and as an open-drain output. As an output, PROCHOT_L is driven low to indicate that the HTC-active state has been entered due to an internal condition, as described by the following text. The minimum assertion and deassertion time for PROCHOT_L is 200 us.

The processor enters the HTC-active state if all of the following conditions are true:

- [D18F3xE8\[HtcCapable\]](#)=1
- [D18F3x64\[HtcEn\]](#)=1
- PWROK=1
- THERMTRIP_L=1
- The processor is not in the C3 ACPI state.

and any of the following conditions are true:

- Tctl is greater than or equal to the HTC temperature limit ([D18F3x64\[HtcTmpLmt\]](#)).
- PROCHOT_L=0

The processor exits the HTC-active state when all of the following are true:

- Tctl is less than the HTC temperature limit ([D18F3x64\[HtcTmpLmt\]](#)).
- Tctl has become less than the HTC temperature limit ([D18F3x64\[HtcTmpLmt\]](#)) minus the HTC hysteresis limit ([D18F3x64\[HtcHystLmt\]](#)) since being greater than or equal to the HTC temperature limit ([D18F3x64\[HtcTmpLmt\]](#)).
- PROCHOT_L=1.

2.10.4.2 Local Hardware Thermal Control (LHTC)

The *LHTC-active state* is characterized by (1) reduced power consumption and (2) reduced performance. While in the LHTC-active state, the processor reduces power consumption by limiting the maximum P-state specified by [D0F0xBC_x1F628\[LhtcActivePstateLimit\]](#). See [2.5.3.1 \[Core P-states\]](#). LHTC is supported only under Battery Power with BAPM disabled. While in the LHTC-active state, software should not change [D0F0xBC_x1F628\[LhtcActivePstateLimit\]](#). Any change to the previous list of fields when in the LHTC-active state can result in undefined behavior.

The LHTC trip point is specified by [D0F0xBC_x1F89C\[LhtcGtempLimitHi\]](#). The LHTC-active state is independent from the HTC-active state. LHTC does not affect PROCHOT_L output and is not affected by PROCHOT_L input.

The processor enters the LHTC-active state if all of the following conditions are true:

- PWROK is asserted.
- The processor is not in the package C6 (PC6) state.
- Tctl is greater than or equal to LHTC temperature limit [D0F0xBC_x1F89C\[LhtcGtempLimitHi\]](#).
- BAPM is disabled. See [2.5.9 \[Bidirectional Application Power Management \(BAPM\)\]](#).

The processor exits the LHTC-active state when the following is true:

- Tctl is less than LHTC temperature low limit ([D0F0xBC_x1F89C\[LhtcGtempLimitLo\]](#)) since being greater than or equal to the LHTC temperature high limit ([D0F0xBC_x1F89C\[LhtcGtempLimitHi\]](#)).

2.10.4.3 Software P-state Limit Control

[D18F3x68 \[Software P-state Limit\]](#) provides a software mechanism to limit the P-state [MSRC001_0061\[CurP-stateLimit\]](#). See [2.5.3 \[CPU Power Management\]](#).

2.10.4.4 THERMTRIP

If the processor supports the THERMTRIP state (as specified by [D18F3xE4 \[Thermtrip Status\]\[ThermtpEn\]](#) or [CPUID Fn8000_0007_EDX\[TTP\]](#), which are the same) and the temperature approaches the point at which the processor may be damaged, the processor enters the THERMTRIP state. The THERMTRIP function is enabled after cold reset (after PWROK asserts and RESET_L deasserts). It remains enabled in all other processor states, except during warm reset (while RESET_L is asserted). The THERMTRIP state is characterized as follows:

- The THERMTRIP_L signal is asserted.
- Nearly all clocks are gated off to reduce dynamic power.
- A low-value VID is generated.
- In addition, the external chipset is expected to place the system into the S5 ACPI state (power off) if THERMTRIP_L is detected to be asserted.

A cold reset is required to exit the THERMTRIP state.

2.11 Root Complex

2.11.1 Overview

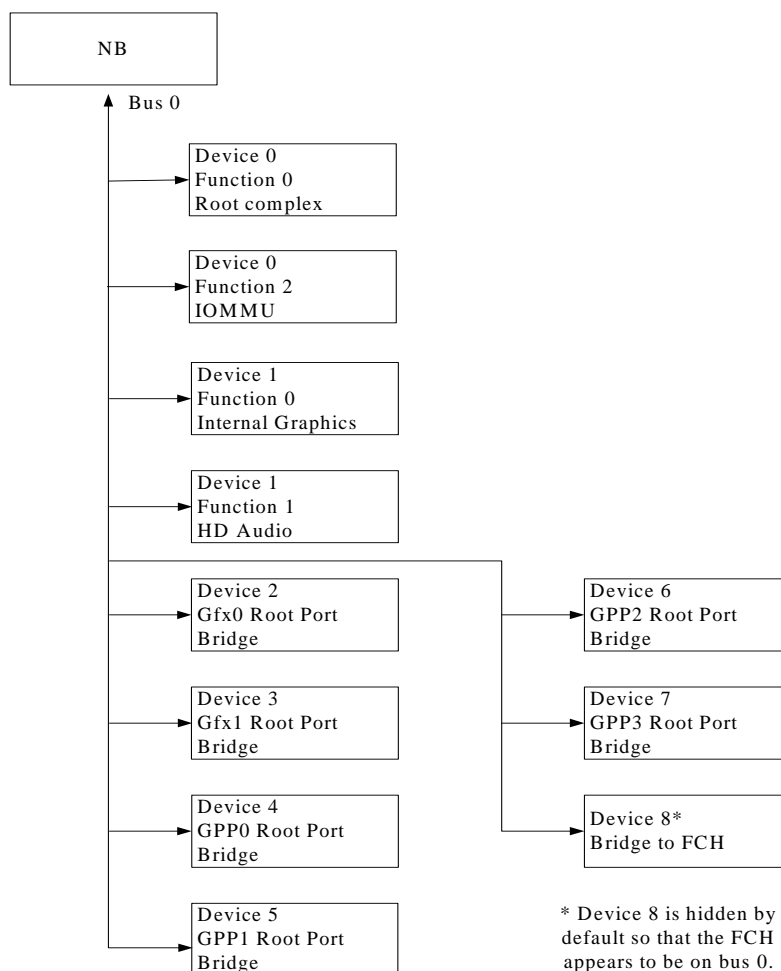


Figure 8: Root complex topology

2.11.2 Interrupt Routing

The RC remaps PCI defined INTx interrupts based on the device number of the virtual bridge or internal device that the interrupt is received from.

Table 47: INTx Mapping

Device Number	INTA	INTB	INTC	INTD
1	INTB	INTC	-	-
2	INTC	INTD	INTA	INTB
3	INTD	INTA	INTB	INTC
4	INTA	INTB	INTC	INTD

Table 47: INTx Mapping

Device Number	INTA	INTB	INTC	INTD
5	INTB	INTC	INTD	INTA
6	INTC	INTD	INTA	INTB
7	INTD	INTA	INTB	INTC

2.11.3 Links

2.11.3.1 Overview

There are 7 configurable ports, which can be divided into 2 groups:

- Gfx: Contains 2 x8 ports. Each port can be limited to lower link widths for applications that require fewer lanes. Additionally, the two ports can be combined to create a single x16 link.
- GPP: Contains 1 x4 UMI and 4 General Purpose Ports (GPP).

Gfx and GPP ports each have a Type 1 Virtual PCI-to-PCI bridge header in the PCI configuration space mapped to devices according to [Figure 8](#).

FS1r2 and FM2 processors can be configured for either Gen1 or Gen2 mode. For FP2 processors see [D18F3x1FC\[Fp2PcieGen2Sup\]](#).

Each PCIe and DDI lane is assigned a unique lane ID that software uses to communicate configuration information to the SMU. [Table 48](#) shows the mappings between lane ID's and lanes.

Table 48: Lane Id Mapping

Lane Id	Lane	Lane Id	Lane	Lane Id	Lane	Lane Id	Lane
0	P_UMI_[T,R]X[P,N]0	10	P_GFX_[T,R]X[P,N]2	20	P_GFX_[T,R]X[P,N]12	30	DP1_TX[P,N]2
1	P_UMI_[T,R]X[P,N]1	11	P_GFX_[T,R]X[P,N]3	21	P_GFX_[T,R]X[P,N]13	31	DP1_TX[P,N]3
2	P_UMI_[T,R]X[P,N]2	12	P_GFX_[T,R]X[P,N]4	22	P_GFX_[T,R]X[P,N]14	32	DP2_TX[P,N]0
3	P_UMI_[T,R]X[P,N]3	13	P_GFX_[T,R]X[P,N]5	23	P_GFX_[T,R]X[P,N]15	33	DP2_TX[P,N]1
4	P_GPP_[T,R]X[P,N]0	14	P_GFX_[T,R]X[P,N]6	24	DP0_TX[P,N]0	34	DP2_TX[P,N]2
5	P_GPP_[T,R]X[P,N]1	15	P_GFX_[T,R]X[P,N]7	25	DP0_TX[P,N]1	35	DP2_TX[P,N]3
6	P_GPP_[T,R]X[P,N]2	16	P_GFX_[T,R]X[P,N]8	26	DP0_TX[P,N]2	36	DP2_TX[P,N]4
7	P_GPP_[T,R]X[P,N]3	17	P_GFX_[T,R]X[P,N]9	27	DP0_TX[P,N]3	37	DP2_TX[P,N]5
8	P_GFX_[T,R]X[P,N]0	18	P_GFX_[T,R]X[P,N]10	28	DP1_TX[P,N]0	38	DP2_TX[P,N]6
9	P_GFX_[T,R]X[P,N]1	19	P_GFX_[T,R]X[P,N]11	29	DP1_TX[P,N]1		

2.11.3.2 Link Configurations

Lanes of the Gfx ports can be assigned to IO links or DDI links.

The following link configurations are supported for the Gfx links:

Table 49: Supported Gfx Port Configurations

D0F0xE4						Gfx Port Lanes ¹			
x0131_0080	x0111_0011	x0211_0011	x0131_8021	x0131_8022	x0131_8013	3:0	7:4	11:8	15:12
0000_0000h	0200_0000h	0200_0000h	7654_3210h	7654_3210h	0000_0001h	x16 Link			
0000_0005h	0001_0000h	0001_0000h	7654_3210h	7654_3210h	0000_0001h	x8 Link		x8 Link	

Table 49: Supported Gfx Port Configurations

DOF0xE4						Gfx Port Lanes ¹			
x0131_0080	x0111_0011	x0211_0011	x0131_8021	x0131_8022	x0131_8013	3:0	7:4	11:8	15:12
0000_0005h	0001_0000h	0000_0300h	7654_3210h	7654_3210h	0000_0001h	x8 Link		x4 Link	DDI
0000_0005h	0001_0000h	0000_0300h	5476_3210h	5476_3210h	0000_0001h	x8 Link		DDI	x4 Link
0000_0000h	0001_0000h	0000_0300h	7654_3210h	7654_3210h	0000_0001h	x8 Link		DDI	DDI
0000_0000h	0001_0000h	0001_0000h	7654_3210h	7654_3210h	0000_0001h	x8 Link		Dual-DVI	
0000_0005h	0000_0300h	0000_0300h	7632_5410h	7632_5410h	0000_0001h	x4 Link	x4 Link	DDI	DDI
0000_0005h	0000_0300h	0001_0000h	7632_5410h	7632_5410h	0000_0001h	x4 Link	x4 Link	Dual-DVI	
0000_0005h	0000_0300h	0001_0000h	7654_3210h	7654_3210h	0000_0004h	x4 Link	DDI	x8 Link	
0000_0005h	0000_0300h	0000_0300h	7654_3210h	7654_3210h	0000_0004h	x4 Link	DDI	x4 Link	DDI
0000_0005h	0000_0300h	0000_0300h	5476_3210h	5476_3210h	0000_0008h	x4 Link	DDI	DDI	x4 Link
0000_0000h	0000_0300h	0000_0300h	7654_3210h	7654_3210h	0000_0001h	x4 Link	DDI	DDI	DDI
0000_0000h	0000_0300h	0001_0000h	7654_3210h	7654_3210h	0000_0001h	x4 Link	DDI	Dual-DVI	
0000_0005h	0000_0300h	0000_0300h	5410_7632h	5410_7632h	0000_0004h	unused	DDI	x4 Link	x4 Link
0000_0000h	0001_0000h	0001_0000h	3210_7654h	3210_7654h	0000_0004h	Dual-DVI		x8 Link	
0000_0005h	0001_0000h	0000_0300h	5410_7632h	5410_7632h	0000_0004h	Dual-DVI		x4 Link	x4 Link
0000_0000h	0001_0000h	0000_0300h	7610_3254h	7610_3254h	0000_0004h	Dual-DVI		x4 Link	DDI
0000_0000h	0001_0000h	0000_0300h	1054_3276h	1054_3276h	0000_0008h	Dual-DVI		DDI	x4 Link
0000_0000h	0001_0000h	0000_0300h	7654_3210h	7654_3210h	0000_0001h	Dual-DVI		DDI	DDI
0000_0000h	0001_0000h	0001_0000h	7654_3210h	7654_3210h	0000_0001h	Dual-DVI		Dual-DVI	

Table 50: Supported Gfx Port Configurations for DOF0xE4_x013[3:1]_804[3:0]

DOF0xE4				Gfx Port Lanes ¹			
x0131_8040	x0131_8041	x0131_8042	x0131_8043	3:0	7:4	11:8	15:12
0000_0000h	0000_0000h	0000_0000h	0000_0000h	x16 Link			
0000_0000h	0000_0000h	0000_0000h	0000_0000h	x8 Link		x8 Link	
0000_0000h	0000_0000h	0000_0000h	0000_0001h	x8 Link		x4 Link	DDI
0000_0000h	0000_0000h	0000_0001h	0000_0000h	x8 Link		DDI	x4 Link
0000_0000h	0000_0000h	0000_0001h	0000_0001h	x8 Link		DDI	DDI
0000_0000h	0000_0000h	0000_0001h	0000_0001h	x8 Link		Dual-DVI	
0000_0000h	0000_0000h	0000_0001h	0000_0001h	x4 Link	x4 Link	DDI	DDI
0000_0000h	0000_0000h	0000_0001h	0000_0001h	x4 Link	x4 Link	Dual-DVI	
0000_0000h	0000_0001h	0000_0000h	0000_0000h	x4 Link	DDI	x8 Link	
0000_0000h	0000_0001h	0000_0000h	0000_0001h	x4 Link	DDI	x4 Link	DDI
0000_0000h	0000_0001h	0000_0001h	0000_0000h	x4 Link	DDI	DDI	x4 Link
0000_0000h	0000_0001h	0000_0001h	0000_0001h	x4 Link	DDI	DDI	DDI
0000_0000h	0000_0001h	0000_0001h	0000_0001h	x4 Link	DDI	Dual-DVI	
0000_0000h	0000_0001h	0000_0000h	0000_0000h	unused	DDI	x4 Link	x4 Link
0000_0001h	0000_0001h	0000_0000h	0000_0000h	Dual-DVI		x8 Link	
0000_0001h	0000_0001h	0000_0000h	0000_0000h	Dual-DVI		x4 Link	x4 Link
0000_0001h	0000_0001h	0000_0000h	0000_0001h	Dual-DVI		x4 Link	DDI
0000_0001h	0000_0001h	0000_0001h	0000_0000h	Dual-DVI		DDI	x4 Link
0000_0001h	0000_0001h	0000_0001h	0000_0001h	Dual-DVI		DDI	DDI

Table 50: Supported Gfx Port Configurations for D0F0xE4_x013[3:1]_804[3:0]

D0F0xE4				Gfx Port Lanes ¹			
x0131_8040	x0131_8041	x0131_8042	x0131_8043	3:0	7:4	11:8	15:12
0000_0001h	0000_0001h	0000_0001h	0000_0001h	Dual-DVI		Dual-DVI	

1. FP2 package processors with TDP <= 25W support 8 lanes only. These can be either lanes [7:0] or [15:8].

To achieve the above configurations, program the following registers:

- Program Gfx function in D0F0xE4_x013[1:0]_0080[StrapBifLinkConfig].
- Program Gfx PIF 0 and Gfx PIF 1 in D0F0xE4_x0[2:1]1[3:0]_0011.
- Program Gfx TX Lane Mux in D0F0xE4_x013[3:0]_8021.
- Program Gfx RX Lane Mux in D0F0xE4_x013[3:0]_8022.
- Program Gfx MasterPciePll in D0F0xE4_x013[3:0]_8013.
- Program OwnSlice in the Gfx registers in D0F0xE4_x013[3:1]_804[3:0][OwnSlice].

The following DP0/DP1 DDI configurations are supported:

Table 51: Supported DP0/DP1 DDI Link Configurations

Lanes[3:0]	Lanes[7:4]
DDI	DDI
DDI (Dual-link DVI)	

The following DP2 DDI configurations are supported:

Table 52: Supported DP2 DDI Link Configurations

Lanes[3:0]	Lanes[6:4]
DDI	Unused
DDI (Dual-link DVI)	

The following link configurations are supported for the GPP links:

Table 53: Supported General Purpose (GPP) Link Configurations

D0F0xE4		GPP Port Lane				
x0130_0080	x0110_0011	Lanes[0:3]	4	5	6	7
0000_0001h	0000_0300h	x4 UMI	x4 Link			
0000_0002h	0000_010Ch	x4 UMI	x2 Link		x2 Link	
0000_0003h	0000_0104h	x4 UMI	x2 Link		x1 Link	x1 Link
0000_0004h	0000_0100h	x4 UMI	x1 Link	x1 Link	x1 Link	x1 Link

2.11.3.3 Clocking

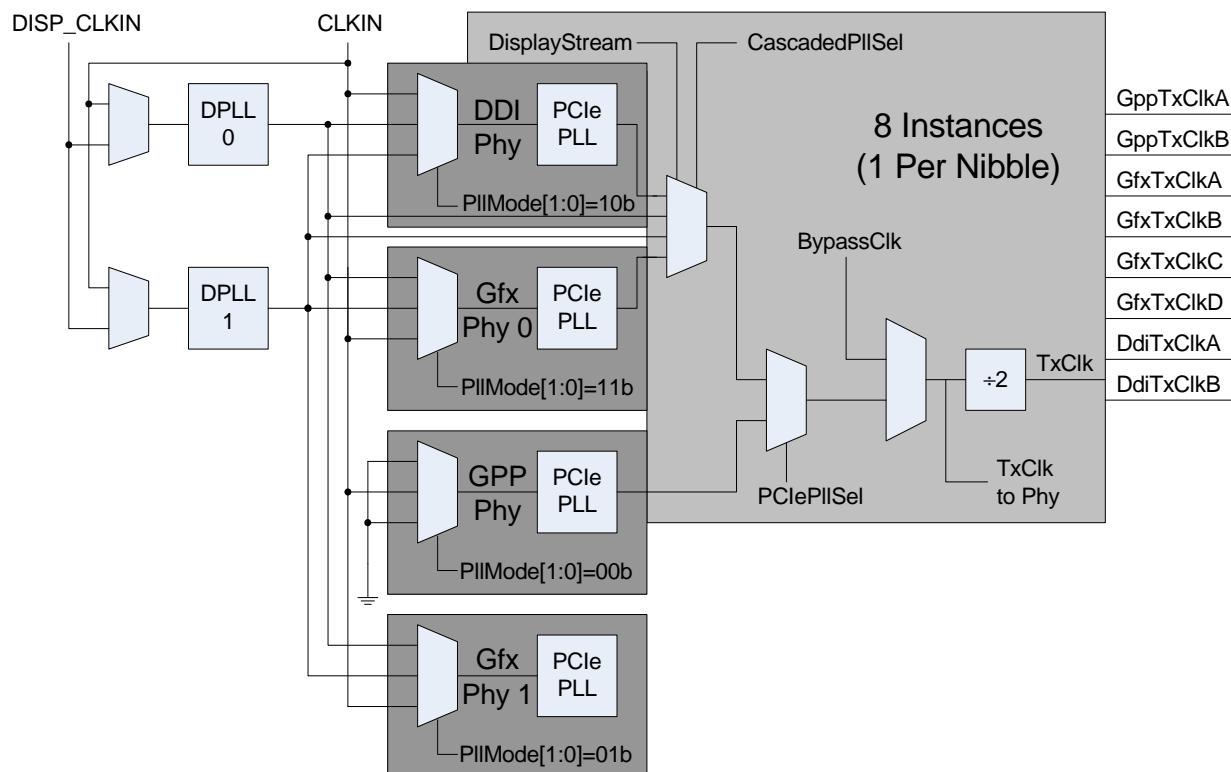


Figure 9: Phy clock configuration

Each phy contains a PLL and 2 clock picker circuits. Each PLL can be controlled independently of the phy and has its clock signal routed to the phys. The PLL in Gfx[8:15] is not connected and should be turned off to conserve power. The GPP PLL is always placed in PCI Express mode and is used as the transmit clock (TxClk) for all PCI Express links. The remaining two PLLs can be configured as clocks for DDI modes.

Setting up the link clocking scheme includes:

- Configuring PLL behavior. See [2.11.4.2.1 \[Clock Configuration\]](#).
- Selecting the clock source for phys, DDI streams, cores, and phy interface.

2.11.4 Root Complex Configuration

2.11.4.1 LPC MMIO Requirements

To ensure proper operation of LPC generated DMA requests, the UMI must be configured to send processor generated MMIO writes that target the LPC bus to the FCH as non-posted writes. To ensure this requirement the MMIO address space of the LPC bus must not be included in the ranges specified by [D18F1x\[1CC:180,BC:80\] \[MMIO Base/Limit\]](#) and non-posted protocol for memory writes must be enabled using the following sequence before LPC DMA transactions are initiated.

1. Configure the FCH to use the non-posted write protocol. See the FCH register specification for configuration details.
2. Program [D0F0xE4_x0101_0010\[UmiNpMemWrite\]](#) = 1.
3. Program [D0F0x98_x06\[UmiNpMemWrEn\]](#) = 1.

2.11.4.2 Link Configuration and Initialization

Link configuration and initialization is performed by the following sequence:

1. IF(UMI) THEN program D0F0x64_x00[NbFchCfgEn] = 1.
2. Clock Configuration (see 2.11.3.3 [Clocking] and 2.11.4.2.1 [Clock Configuration]).
3. 2.11.4.2.2 [Link Configuration and Core Initialization]
4. 2.11.4.2.3 [Link Training]
5. 2.11.4.4 [Power Management]
6. Lock link configuration registers.
 - Program D0F0xE4_x0[2:1]01_0010[HwInitWrLock] = 1.
 - Program D0F0x64_x00[HwInitWrLock] = 1.
7. IF(UMI) THEN program D0F0x64_x00[NbFchCfgEn] = 0.

2.11.4.2.1 Clock Configuration

Clock configuration is required after each cold or warm reset performed. The configuration applies to all phys, regardless of being configured as an IO link or DDI link.

1. Program the PLL to be powered off:
 - A. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllPowerStateInOff]=111b.
 - B. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllRampUpTime]=000b.
 - C. Program D0F0xE4_x0[2:1]1[3:0]_0010[LS2ExitTime]=000b.
2. Power down sub-link PLLs:
 - A. Program D0F0xE4_x0130_8023[LaneEnable]=0Fh. See D0F0xE4_x013[3:0]_8023
 - B. Program D0F0xE4_x0131_8023[LaneEnable]=00h.
 - C. Program D0F0xE4_x0132_8023[LaneEnable]=00h.
3. Wait for D0F0xE4_x0[2:1]1[3:0]_0015[7:0]==FFh.
4. Program the PLL mode:
 - A. Program D0F0xE4_x0121_2005[PllMode]=01b. See D0F0xE4_x0[2:1]2[3:0]_2005
 - B. Program D0F0xE4_x0221_2005[PllMode] = 00b.
 - C. Program D0F0xE4_x0122_2005[PllMode] = 10b.
5. Enable the individual lanes:
 - A. Program D0F0xE4_x0130_8023[LaneEnable]=FFh. See D0F0xE4_x013[3:0]_8023
 - B. Program D0F0xE4_x0131_8023[LaneEnable]=FFFFh.
 - C. Program D0F0xE4_x0132_8023[LaneEnable]=FFh.
 - D. Wait for D0F0xE4_x0[2:1]1[3:0]_0015[7:0]==FFh.
6. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllPowerStateInOff]=000b.

2.11.4.2.2 Link Configuration and Core Initialization

Link configuration is done on a per link basis. Lane reversal, IO link/DDI link selection, and lane enablement is configured through this sequence.

1. Place software-reset module into blocking mode:
 - A. Program D0F0xE4_x013[3:0]_8062[ConfigXferMode]=0.
 - B. Program D0F0xE4_x013[3:0]_8062[BlockOnIdle]=0.
2. If the link is an IO link, Program D0F0xE4_x0[2:1]01_0011[DynamicClkLatency]=Fh.
3. Program D0F0xE4_x013[1:0]_0080 per Table 49 and Table 53.
4. Program D0F0xE4_x013[3:0]_8021 per Table 49 and Table 53.
5. Program D0F0xE4_x0[2:1]1[3:0]_0010[RxDetectTxPwrMode]=1.
6. Program D0F0xE4_x0[2:1]1[3:0]_0010[LS2ExitTime]=000b.

7. Program D0F0xE4_x013[3:0]_8013[MasterPciePllA, MasterPciePllB, MasterPciePllC, MasterPciePllD] per Table 49.
8. Initiate core reconfiguration sequence:
 - A. Program D0F0xE4_x013[3:0]_8062[ReconfigureEn]=1.
 - B. Program D0F0xE4_x013[3:0]_8060[Reconfigure]=1.
 - C. Wait for D0F0xE4_x013[3:0]_8060[Reconfigure]==0.
 - D. Program D0F0xE4_x013[3:0]_8062[ReconfigureEn]=0.
9. Return software-reset module to non-blocking mode:
 - A. Program D0F0xE4_x013[3:0]_8062[ConfigXferMode]=1.
10. Program D[8:2]F0xE4_xC1[StrapReverseLanes] if necessary.
11. Program D0F0xE4_x0[2:1]1[3:0]_0011 per Table 49 and Table 53.
12. Program D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5 per Table 105.
13. For each link mapped to DDI:
 - A. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllPowerStateInTxs2]=111b.
 - B. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllPowerStateInOff]=111b.
 - C. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllRampUpTime]=010b.
14. For each nibble that has no PCIe lanes in use:
 - A. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllPowerStateInOff]=111b.
 - B. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][PllPowerStateInTxs2]=111b.
 - C. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][TxPowerStateInTxs2]=111b.
 - D. Program D0F0xE4_x0[2:1]1[3:0]_001[3:2][RxPowerStateInRxs2]=111b.
15. For each lane that is not in use, program the corresponding D0F0xE4_x013[3:0]_8023[LaneEnable]=0.
16. If the link is a DDI link:
 - A. Program D0F0xE4_x013[3:1]_804[3:0][OwnSlice] per Table 49.
17. Configure PIF parings and disable ganged mode for UMI:
 - A. Program D0F0xE4_x0110_0011=0000_0300h.
 - B. Program D0F0xE4_x0120_6[3:2][8,0]5[GangedModeEn]=0.

2.11.4.2.3 Link Training

Link training is performed on a per link basis. BIOS may train the links in parallel.

2.11.4.3 Miscellaneous Features

2.11.4.3.1 Straps

1. Program D0F0xE4_x013[3:0]_8011[StrapBifValid]=1.
2. Program strap values.
3. Program D0F0xE4_x013[3:0]_8011[StrapBifValid]=0.

2.11.4.3.2 Lane Reversal

Normally, the lanes of each port are physically numbered from n-1 to 0 where n is the number of lanes assigned to the port. Physical lane numbering can be reversed according to the following methods:

- To reverse the physical lane numbering for a specific port, program D[8:2]F0xE4_xC1[StrapReverseLanes]=1 according to the sequence in 2.11.4.3.1 [Straps].
- To reverse the physical lane numbering for all ports in the GPP or GFX interfaces, program D0F0xE4_x0[2:1]01_00C0[StrapReverseAll]=1 according to the sequence in 2.11.4.3.1 [Straps].

Note that logical port numbering is established during link training regardless of the physical lane numbering.

2.11.4.3.3 Link Speed Changes

Link speed changes can only occur on Gen2 capable links. To verify that Gen2 speeds are supported verify `D[8:2]F0x64[LinkSpeed]==02h`.

2.11.4.3.3.1 Software Initiated Link Speed Changes

The following programming sequence describes a software initiated speed change from Gen1 to Gen2:

1. Verify `D[8:2]F0xE4_xA4[LcGen2EnStrap]==1`. If `LcGen2EnStrap` is not 1, program it to 1.
2. Program `D[8:2]F0x88[TargetLinkSpeed]=2`.
3. Program `D[8:2]F0xE4_xA4[LcGoToRecovery]=1`.
4. Verify `D[8:2]F0xE4_xA4[LcOtherSideSupportsGen2]==1`, otherwise stop.
5. If `D[8:2]F0x68[LinkSpeed]==2h`, stop. The link is already at Gen2 speed.
6. If `D[8:2]F0xE4_xA4[LcSpeedChangeAttemptFailed]==1`, stop. The maximum number of speed negotiation failures have been reached.
7. Ensure `D[8:2]F0xE4_xA4[LcForceDisSwSpeedChange]=0`.
8. Ensure `D[8:2]F0x88[HwAutonomousSpeedDisable]=0`.
9. Program `D[8:2]F0xE4_xA4[LcInitiateLinkSpeedChange]=1`.

The following programming sequence describes a software initiated speed change from Gen2 to Gen1:

1. Program `D[8:2]F0x88[TargetLinkSpeed]=1h`.
2. If `D[8:2]F0x68[LinkSpeed]==1h`, stop. The link is already at Gen1 speed.
3. If `D[8:2]F0xE4_xA4[LcSpeedChangeAttemptFailed]==1`, stop. The maximum number of speed negotiation failures have been reached.
4. Ensure `D[8:2]F0xE4_xA4[LcForceDisSwSpeedChange]=0`.
5. Ensure `D[8:2]F0x88[HwAutonomousSpeedDisable]=0`.
6. Program `D[8:2]F0xE4_xA4[LcInitiateLinkSpeedChange]=1`.

2.11.4.3.3.2 Autonomous Link Speed Changes

To enable autonomous speed changes on a per port basis:

1. Program `D[8:2]F0x88[TargetLinkSpeed]=2h`.
2. Program `D0F0xE4_x013[1:0]_0[C:8]03[StrapBifDeemphasisSel]=1`.
3. Program `D[8:2]F0xE4_xA4[LcGen2EnStrap]=1`.
4. Program `D[8:2]F0xE4_xC0[StrapAutoRcSpeedNegotiationDis]=0`.
5. Program `D[8:2]F0xE4_xA4[LcMultUpstreamAutoSpdChangEn]=1`.
6. Program `D[8:2]F0xE4_xA2[LcUpconfigureDis]=0`.

2.11.4.3.4 Deemphasis

Deemphasis strength can be changed on a per-port basis by programming `D[8:2]F0xE4_xB5[LcSelectDeemphasis]`.

2.11.4.4 Power Management

2.11.4.4.1 Link States

To enable support for L1 program `D[8:2]F0xE4_xA0[LcL1Inactivity]=6h`.

To enable support for L0s:

- Program `D[8:2]F0xE4_xA1[LcDontGotoL0sifL1Armed]=1`.
- Program `D[8:2]F0xE4_xA0[LcL0sInactivity]=9h`.

2.11.4.4.2 Dynamic Link-width Control

Dynamic link-width control is a power saving feature that reconfigures the link to run with fewer lanes. The inactive lanes are turned off to conserve power.

The GFX links can switch among widths of: x1, x2, x4, x8, and x16.

The GPP links can switch among widths of: x1, x2, and x4.

There are 3 link-width control mechanisms:

- Long Reconfiguration: The link is brought down and retrained to a different width. This mechanism is only supported on links between AMD products.
- Short Reconfiguration: The link is retrained to a different width by going through the recovery state of the LTSSM. This mechanism is only supported on links between AMD products.
- Up/Down Reconfiguration: The link is retrained according to the PCI Express specification. This mechanism is only available between Gen2 devices.

The core has the capability to turn off the inactive lanes of trained links. To enable this feature program [D\[8:2\]F0xE4_xA2\[LcDynLanesPwrState\]=11b](#).

2.11.4.5 Link Test and Debug Features

2.11.4.5.1 Compliance Mode

To enable Gen1 software compliance mode program [D\[8:2\]F0xE4_xC0\[StrapForceCompliance\]=1](#) for each port to be placed in compliance mode.

To enable Gen2 software compliance mode:

1. BIOS enables Gen2 capability by programming [D0F0xE4_x0\[2:1\]01_00C1\[StrapGen2Compliance\]=1](#).
2. Program [D0F0xE4_x013\[1:0\]_0\[C:8\]03\[StrapBifDeemphasisSel\]=1](#) for each port to be placed in compliance mode.
3. Program [D\[8:2\]F0x88\[TargetLinkSpeed\]=2h](#) for each port to be placed in compliance mode.
4. Program [D\[8:2\]F0x88\[EnterCompliance\]=1](#) for each port to be placed in compliance mode.

2.11.5 BIOS Timer

The root complex implements a 32-bit microsecond timer (see [D0F0xE4_x0130_80F0](#) and [D0F0xE4_x0130_80F1](#)) that the BIOS can use to accurately time wait operations between initialization steps. To ensure that BIOS waits a minimum number of microseconds between steps BIOS should always wait for one microsecond more than the required minimum wait time.

2.12 IOMMU

The processor includes an IOMMU revision 2 for FM2 packages only. See the *AMD I/O Virtualization Technology (IOMMU) Specification*.

2.12.1 IOMMU Configuration Space

The IOMMU configuration space consists of the following four groups:

- PCI Configuration space. See [3.4 \[Device 0 Function 2 \(IOMMU\) Configuration Registers\]](#).
- IOMMU Memory Mapped Register space. See [3.15 \[IOMMU Memory Mapped Registers\]](#).

- IOMMU L1 Indexed space accessed through [D0F2xF8 \[IOMMU L1 Config Index\]](#).
- IOMMU L2 Indexed space accessed through [D0F2xF0 \[IOMMU L2 Config Index\]](#).

2.12.2 IOMMU Initialization

BIOS should perform the following steps to initialize the IOMMU:

1. Program [D0F0x64_x0D\[Dev0Fn2RegEn\]](#)=1.
2. Program [D0F2x44 \[IOMMU Base Address Low\]](#) and [D0F2x48 \[IOMMU Base Address High\]](#) to allocate a 512K region of MMIO space for IOMMU memory mapped registers. This region of MMIO space is reserved for IOMMU and BIOS must not allocate it for use by system software.
3. Program [D0F2x44\[IommuEnable\]](#)=1.
4. Program [D\[8:2\]F0xE4_xC1\[StrapExtendedFmtSupported\]](#)=1 and [D\[8:2\]F0xE4_xC1\[StrapE2EPrefixEn\]](#)=1.
5. Program the registers with BIOS recommendations in L1 ([D0F2xF8](#)) and L2 ([D0F2xF0](#)) indexed space.
6. Check if any PCIe devices in the system support the Phantom Function. For each PCIe core that has a connected device advertising support for the Phantom Function, program [D0F2xFC_x07_L1sel\[3:0\]\[PhantomSupEn\]](#)=1 for the L1 corresponding to that PCIe core.
7. If a PCIe port is hot-plug capable, then program [D0F2xFC_x07_L1sel\[3:0\]\[PhantomSupEn\]](#)=1 for the L1 corresponding to the PCIe core.
8. If at least one PCIe to PCI-x bridge exists on a PCIe port or a HotPlug capable PCIe slot is present on a PCIe port then program [D0F2xFC_x0D_L1sel\[3:0\]\[VOQPortBits\]](#)=111b for the L1 corresponding to the particular PCIE core.

2.13 System Management Unit (SMU)

The system management unit (SMU) is a subcomponent of the northbridge that is responsible for a variety of system and power management tasks during boot and runtime. The SMU contains a microcontroller to assist with many of these tasks.

2.13.1 Software Interrupts

The microcontroller can be interrupted to cause it to perform several initialization and runtime tasks. BIOS and ACPI methods can interrupt the SMU to request a specific action using the following sequence:

1. Wait for `D0F0xBC_xE0003004[IntDone]==1`.
2. Program `D0F0xBC_xE0003000[ServiceIndex]` to the desired service index and toggle `D0F0xBC_xE0003000[IntReq]`. This may be done in a single write.
3. Wait for `D0F0xBC_xE0003004[IntAck]==1`.

After performing the steps above, software may continue execution before the interrupt has been serviced. However, software should not rely on the results of the interrupt until the service is complete (see `D0F0xBC_xE0003004[IntDone]`). Interrupting the SMU with a service index that does not exist results in undefined behavior.

Table 54: BIOS Services

Service Index	Notes
16h	Description: <code>SMC_MSG_CONFIG_NBDPM</code> . Enables NB P-state Adjustments. See 2.5.4.1 [NB P-states].
	Input: <code>D0F0xBC_x1F5F8</code>
	Output: <code>D0F0xBC_x1F5FC</code>
	Firmware revision:

2.14 Graphics Processor (GPU)

The processor contains an integrated DX11 compliant graphics processor.

2.14.1 GPU PCI Interface

BIOS must configure the PCI interface block of the GPU before the GPU can be accessed. The PCI interface block is configured by programming the `D0F0x64_x1C` [Internal Graphics PCI Control 1] register. Two writes to this register are required. The first write sets the configuration and must have `D0F0x64_x1C[WriteDis]=0`. The second write uses the same write data with `D0F0x64_x1C[WriteDis]=1`.

2.14.2 Graphics Memory Controller (GMC)

The graphics memory controller is responsible for servicing memory requests from the different blocks within the GPU and forwarding routing them to the appropriate interface. The GMC is also responsible for translating GPU virtual address to GPU physical addresses and for translating GPU physical addresses to system addresses.

2.14.3 Frame Buffer (FB)

The frame buffer is defined as the portion of system memory dedicated for GPU use.

Table 55: Recommended Frame Buffer Configurations

System Memory Size	Frame Buffer Size
< 1GB	64 MB
>= 1GB & < 2GB	256 MB
>= 2GB & < 4GB	512 MB
>= 4GB	768 MB

2.15 RAS Features

This section applies reliability, availability, and serviceability, or RAS, and related considerations.

2.15.1 Machine Check Architecture

The processor contains logic and registers to control detection, corrective action, logging, and reporting of errors in the data or control paths in each core and the northbridge.

This section assumes familiarity with the *AMD64 Architecture Programmer's Manual Volume 2: System Programming* chapter titled "Machine Check Mechanism". See [1.2 \[Reference Documents\]](#).

The ability of hardware to generate a machine check exception upon an error is indicated by [CPUID Fn0000_0001_EDX\[MCE\]](#).

2.15.1.1 Machine Check Registers

[CPUID Fn0000_0001_EDX\[MCA\]](#) indicates the presence of the following machine check registers:

- [MSR0000_0179 \[Global Machine Check Capabilities \(MCG_CAP\)\]](#)
 - Reports how many machine check register banks are supported.
- [MSR0000_017A \[Global Machine Check Status \(MCG_STAT\)\]](#)
- [MSR0000_017B \[Global Machine Check Exception Reporting Control \(MCG_CTL\)\]](#)

The error reporting register banks supported are:

- **MC0**: load-store unit (LS), including data cache.
- **MC1**: instruction fetch unit (IF), including instruction cache.
- **MC2**: combined unit (CU), including L2 cache.
- **MC3**: Reserved.
- **MC4**: northbridge (NB). These MSR's are also accessible from configuration space. There is only one NB error reporting bank, independent of the number of cores.
- **MC5**: execution unit (EX), including mapper/scheduler/retire/execute functions and fixed-issue reorder buffer.
- **MC6**: floating point unit (FP).

The register types within each bank are:

- **MCi_CTL**: The Machine Check Control Register: Enables error reporting.
- **MCi_STATUS**: The Machine Check Status Register: Logs information associated with errors.
- **MCi_ADDR**: The Machine Check Address Register: Logs address information associated with errors.
- **MCi_MISC**: The Machine Check Miscellaneous Register: Log miscellaneous information associated with errors, as defined by each error type.
- **MCi_CTL_MASK**: The Machine Check Control Mask Register: Inhibit detection of an error source unless otherwise specified.

[Table 56](#) identifies the registers associated with each register bank:

Table 56: MCA Register Cross-Reference Table

Register Bank	MCA Register				
	CTL	STATUS	ADDR	MISC	CTL_MASK
MC0	MSR0000_0400	MSR0000_0401	MSR0000_0402	MSR0000_0403	MSRC001_0044
MC1	MSR0000_0404	MSR0000_0405	MSR0000_0406	MSR0000_0407	MSRC001_0045
MC2	MSR0000_0408	MSR0000_0409	MSR0000_040A	MSR0000_040B	MSRC001_0046
MC3	MSR0000_040C	MSR0000_040D	MSR0000_040E	MSR0000_040F	MSRC001_0047
MC4	MSR0000_0410	MSR0000_0411	MSR0000_0412	MSR0000_0413 MSRC000_0408	MSRC001_0048
MC5	MSR0000_0414	MSR0000_0415	MSR0000_0416	MSR0000_0417	MSRC001_0049
MC6	MSR0000_0418	MSR0000_0419	MSR0000_041A	MSR0000_041B	MSRC001_004A

Uncorrectable errors that are enabled in MCi_CTL result in reporting to software via machine check exceptions. Some errors increment a counter in MCi_MISC , which may trigger an interrupt (see 2.15.1.7 [Error Thresholding]).

Each register bank i implements a number of machine check miscellaneous registers, denoted as MCi_MISCj , where j goes from 0 to a maximum of 8. The presence of valid information in the first MISC register in the bank (MCi_MISC0) is indicated by $MCi_STATUS[MiscV]$, and in subsequent MISC registers by $MCi_MISCj[Val]$ in the target register. If there is more than one MISC register in a given bank, a non-zero value in $MCi_MISC0[BlkPtr]$ points to the contiguous block of additional registers.

2.15.1.2 Machine Check Errors Classes

The classes of machine check errors are:

- Uncorrectable
- Deferred
- Correctable

Uncorrectable errors cannot be corrected by hardware or microcode and may have caused the loss of data or corruption of processor state. Uncorrectable errors, if not masked from logging, update the status and address registers, and if enabled for reporting, cause a machine check exception. If there is information in the status and address registers from a previously logged deferred or correctable error, it is overwritten. If an uncorrectable error is masked from logging, the error is ignored by hardware (exceptions are noted in the register definitions). If an uncorrectable error is disabled from reporting, containment of the error and logging/reporting of subsequent errors may be affected. Therefore, unmasked uncorrectable errors should be enabled for reporting for normal operation. Uncorrectable errors should only be disabled from reporting for debug purposes.

Deferred errors are errors that cannot be corrected by hardware, but that do not cause an immediate interruption in program flow, loss of data integrity, or corruption of processor state. They are errors that indicate that data has been corrupted but contained; no exception is generated because the data has not been referenced by a core or an IO link. If deferred errors are enabled for logging, the status and address registers in the corresponding register bank are written with information that identifies the source of the error. If there is information in the status and address registers from a previously logged correctable error, it is overwritten.

Correctable errors are corrected by hardware or microcode and cause no loss of data or corruption of processor state (unless disabled by implementation-specific bits in the control registers for test or debug reasons). If they are enabled for logging, the status and address registers in the corresponding register bank are written with

information that identifies the source of the error. Correctable errors are not reported via machine check exceptions; some correctable errors may be reported via error thresholding (see 2.15.1.7 [Error Thresholding]).

The implications of these categories of errors are:

1. Uncorrected error; the problem was not dealt with by hardware.
 - Operationally (error handling), action does need to be taken, because program flow is affected.
 - Diagnostically (fault management), software may collect information to determine if any components should be de-configured or serviced.
 - Examples include:
 - Uncorrectable ECC, no way to avoid passing it to process.
 - Poison data consumed (as opposed to created), no way to avoid passing it to process or link.
2. Deferred error; the problem was partly dealt with by containment.
 - Operationally, no immediate action needs to be taken, because program flow has not been affected. However, steps may be taken by software to prevent access to the data in error.
 - Diagnostically, software may collect information to determine if any components should be de-configured or serviced.
 - Examples include:
 - Uncorrectable ECC, converted to poison data.
3. Corrected error; the problem was dealt with by hardware.
 - Operationally, no action needs to be taken, because program flow is unaffected.
 - Diagnostically, software may collect information to determine if any components should be de-configured or serviced.
 - Examples include:
 - Correctable ECC, corrected in-line.

For debug observability, [D18F3x180](#)[ChgUcToCeEn] can be used to convert NB uncorrectable errors to correctable errors.

Machine check conditions can be simulated by using [MSRC001_0015](#)[McStatusWrEn]. This is useful for debugging machine check handlers. See 2.15.2 [Error Injection and Simulation] for more detail.

2.15.1.3 Error Detection, Action, Logging, and Reporting

Error detection:

- Core error detection is enabled if not masked by `MCi_CTL_MASK` ([MSRC001_0044-MSRC001_0046](#), [MSRC001_0049-MSRC001_004A](#)).
- NB error detection (MC4) is not affected by `MC4_CTL_MASK` ([MSRC001_0048](#)).
- Error masking is performed regardless of MCA bank enablement in `MCG_CTL` ([MSR0000_017B](#)).

Error actions are enabled if all of the following are true:

- Error detection is enabled.
- The condition that enables the action to be taken for a detected error varies. The EAC (Error Action Condition) column of the error description tables defines whether `MCG_CTL[i]` affects whether the error action is taken.
 - D: Detected. The error action is taken if the error is detected. `MCG_CTL[i]` does not affect if the error action is taken.
 - E: Enabled. The error action is taken only if the bank is enabled by `MCG_CTL[i]==1`.

Error logging is enabled if all of the following are true:

- Error detection is enabled.

- The MCA bank is enabled in MCG_CTL ([MSR0000_017B](#)).
- MC4/NB only: Error logging is not masked. See MC4_CTL_MASK ([MSRC001_0048](#)).

Error reporting is enabled if all of the following are true:

- Error logging is enabled.
- The corresponding enable bit for the error in MCI_CTL is set to 1.

A machine check exception will be generated if all the following are true:

- The error is enabled for reporting.
- The error is uncorrectable.
- CR4.MCE is enabled.

Notes:

1. If error reporting is enabled but CR4.MCE is disabled or the machine check handler cannot be invoked for other reasons, then a reportable error will cause the system to enter shutdown.
2. If error reporting is disabled, the setting of CR4.MCE has no effect.
3. If an uncorrectable error is disabled from reporting, containment of the error and logging/reporting of subsequent errors may be affected. Therefore, unmasked uncorrectable errors should be enabled for reporting for normal operation. Uncorrectable errors should only be disabled from reporting for debug purposes.
4. Errors not associated with a specific core are reflected to core 0 of the compute unit. The error description tables identify which errors are associated or not associated with a specific core of the compute unit.

Throughout the MCA register descriptions, the terms “enabled” and “disabled” generally refer to reporting, and the terms “masked” and “unmasked” generally refer to logging, unless otherwise noted.

Some logged errors increment a counter in MCI_MISC, which may trigger an interrupt (see [2.15.1.7 \[Error Thresholding\]](#)).

2.15.1.3.1 MCA conditions that cause Shutdown

The following architectural conditions cause the processor to enter the Shutdown state; see “Machine-Check Errors” and subsections in APM volume 2 for more detail; see [1.2 \[Reference Documents\]](#):

- Attempting to generate an MCE when CR4.MCE=0.
- Attempting to generate an MCE when [MSR0000_017A\[MCIP\]=1](#).

The following non-architectural conditions cause the processor to enter the Shutdown state:

- EX “Retire dispatch queue parity” error. See [Table 230 \[EX Error Descriptions\]](#).
- EX “Mapper checkpoint array parity” error if UC=1. See [Table 230 \[EX Error Descriptions\]](#).

2.15.1.3.2 Error Logging During Overflow

An error to be logged when the status register contains valid data can result in an overflow condition. During error overflow conditions, the new error may not be logged or an error which has already been logged in the status register may be overwritten. For the rules on error overflow, priority, and overwriting, see [MSR0000_0401\[Overflow\]](#) and [MSR0000_0411\[Overflow\]](#).

Uncorrectable errors require software intervention. Therefore, when an uncorrectable error cannot be logged, critical error information may have been lost, and MCI_STATUS[PCC] may be set. If PCC is indicated, software should terminate system processing to prevent data corruption (see [2.15.1.6 \[Handling Machine Check Exceptions\]](#)). If PCC is not indicated, any MCA data lost due to overflow was informational only and not critical to system hardware operation.

Table 57 indicates which errors are overwritten in the error status registers.

Table 57: Overwrite Priorities for All Banks

			Older Error			
			Uncorrectable		Correctable	
			Enabled	Disabled	Enabled	Disabled
Newer Error	Uncorrectable	Enabled	-	-	Overwrite	Overwrite
		Disabled	-	-	Overwrite	Overwrite
	Correctable	Enabled	-	-	-	-
		Disabled	-	-	-	-

2.15.1.4 MCA Initialization

The following initialization sequence must be followed:

- MCI_CTL_MASK registers ([MSRC001_0044](#) to [MSRC001_004A](#)):
 - BIOS must initialize the mask registers to inhibit error detection prior to the initialization of MCI_CTL and [MSR0000_017B](#).
 - BIOS must not clear MASK bits that are reset to 1.
- The MCI_CTL registers must be initialized prior to enabling the error reporting banks in MCG_CTL.

If initializing after a cold reset (see [D18F0x6C](#)[ColdRstDet]), then BIOS must clear the MCI_STATUS MSRs (see [Table 56](#)). If initializing after a warm reset, then BIOS should check for valid MCA errors and if present save the status for later use (see [2.15.1.6 \[Handling Machine Check Exceptions\]](#)).

BIOS that wishes to ensure continued operation in the event that a machine check occurs during boot may write MCG_CTL with all ones and write zeros into each MCI_CTL. With these settings, a machine check error will result in MCI_STATUS being written without generating a machine check exception or a system shutdown. BIOS may then poll MCI_STATUS during critical sections of boot to ensure system integrity. Before passing control to the operating system, BIOS should restore the values of those registers to what the operating system is expecting. (Note that using MCI_CTL to disable error reporting on uncorrectable errors will affect error containment; see [2.15.1.3 \[Error Detection, Action, Logging, and Reporting\]](#).) Alternatively, the BIOS initialize MCA without setting CR4.MCE; This will result in a system shutdown on any machine check which would have caused a machine check exception (followed by a reboot if configured in the chipset).

2.15.1.5 Error Code

The MCI_STATUS[ErrorCode] field contains information on the logged error. [Table 58 \[Error Code Types\]](#) identifies how to decode ErrorCode. The MCI_STATUS[ErrorCodeExt] field contains detailed, model-specific information that is used for error diagnosis but not error handling; see [2.15.1.6 \[Handling Machine Check Exceptions\]](#).

For a given error reporting bank, Error Code is used in conjunction with other MCI_STATUS fields to identify the Error Type. Details for each Error Type are described in the error signatures tables accompanying the MCI_STATUS register for each bank:

- MC0; [Table 213](#).
- MC1; [Table 216](#).
- MC2; [Table 220](#).
- MC3; Reserved.
- MC4; [Table 223](#) and [Table 224](#).

- MC5; [Table 231](#).
- MC6; [Table 234](#).

Table 58: Error Code Types

Error Code	Error Code Type	Description
0000_0000_0001_TTLL	TLB	Errors in the GART TLB cache. TT = Transaction Type LL = Cache Level
0000_0001_RRRR_TTLL	MEM	Errors in the cache hierarchy (not in NB) RRRR = Memory Transaction Type TT = Transaction Type LL = Cache Level
0000_1PPT_RRRR_IILL	BUS	General bus errors including link and DRAM PP = Participation Processor T = Timeout RRRR = Memory Transaction Type II = Memory or IO LL = Cache Level

Table 59: Error Codes: Transaction Type

TT	Transaction Type
00	I: Instruction
01	D: Data
10	G: Generic
11	Reserved

Table 60: Error codes: cache level

LL	Cache Level
00	Reserved
01	L1: Level 1
10	L2: Level 2
11	LG: Generic

Table 61: Error Codes: Memory Transaction Type

RRRR	Memory Transaction Type
0000	GEN: Generic. Includes scrub errors.
0001	RD: Generic Read
0010	WR: Generic Write
0011	DRD: Data Read
0100	DWR: Data Write
0101	IRD: Instruction Fetch

Table 61: Error Codes: Memory Transaction Type

RRRR	Memory Transaction Type
0110	Prefetch
0111	Evict
1000	Probe (Snoop)

Table 62: Error Codes: Participation Processor

PP	Participation Processor
00	SRC: Local node originated the request
01	RES: Local node responded to the request
10	OBS: Local node observed the error as a third party
11	GEN: Generic

Table 63: Error Codes: Memory or IO

II	Memory or IO
00	MEM: Memory Access
01	Reserved
10	IO: IO Access
11	GEN: Generic

2.15.1.6 Handling Machine Check Exceptions

A machine check handler is invoked to handle an exception for a particular core. Because MCA registers are generally not shared among cores, the handler does not need to coordinate register usage with handler instances on other cores. (Those few MCA registers which are shared are noted in the register description. See also [2.4.1.1 \[Registers Shared by Cores in a Compute Unit\]](#).) For access to the NB MCA registers, [D18F3x44\[NbMcaToMstCpuEn\]](#) allows a single core (the NBC) to access the registers through MSR space without contention from other cores. This organization of registers on a per core basis allows independent execution, simplifies exception handling, and reduces the number of conditions which are globally fatal.

At a minimum, the machine check handler must be capable of logging error information for later examination. The handler should log as much information as is needed to diagnose the error.

More thorough exception handler implementations can analyze errors to determine if each error is recoverable. If a recoverable error is identified, the exception handler can attempt to correct the error and restart the interrupted program. An error may not be recoverable for the process or virtual machine it directly affects, but may be containable, so that other processes or virtual machines in the system are unaffected and system operation is recovered; see [2.15.1.6.1 \[MCA Differentiation Between System-Fatal and Process-Fatal Errors\]](#).

Machine check exception handlers that attempt to recover must be thorough in their analysis and the corrective actions they take. The following guidelines should be used when writing such a handler:

- Data collection:
 - All status registers in the error reporting banks must be examined to identify the cause of the machine check exception.

- Read [MSR0000_0179](#)[Count] to determine the number of status registers visible to the core. The status registers are numbered from 0 to one less than the value found in [MSR0000_0179](#)[Count]. For example, if the Count field indicates five status registers are supported, they are numbered MC0_STATUS to MC4_STATUS. These are generically referred to as MC_i_STATUS.
- Check the valid bit in each status register (MC_i_STATUS[Val]). The remainder of the status register does not need to be examined when its valid bit is clear.
- When identifying the error condition and determining how to handle the error, portable exception handlers should examine the following MC_i_STATUS fields: ErrorCode, UC, PCC, CECC, UECC, Deferred, Poison. The expected settings of these and other fields in MC_i_STATUS are identified in the error signatures tables which accompany the descriptions of each MCA status register. See [2.15.1.5 \[Error Code\]](#) for a discussion of error codes and pointers to the error signatures tables.
 - MC_i_STATUS[ErrorCodeExt] should generally not be used by portable code to identify the error condition because it is model specific. ErrorCodeExt is useful in determining the error subtype for root cause analysis.
- Error handlers should collect all available MCA information (status register, address register, miscellaneous register, etc.), but should only interrogate details to the level which affects their actions. Lower level details may be useful for diagnosis and root cause analysis, but not for error handling.
- Recovery:
 - Check the valid MC_i_STATUS registers to see if error recovery is possible. Error recovery is not possible when the processor context corrupt indicator (MC_i_STATUS[PCC]) is set to 1.
 - The error overflow status indicator (MC_i_STATUS[Overflow]) does not indicate whether error recovery is possible.
 If error recovery is not possible, the handler should log the error information and return to the operating system for system termination.
 - Check MC_i_STATUS[UC] to see if the processor corrected the error. If UC is set, the processor did not correct the error, and the exception handler must correct the error prior to attempting to restart the interrupted program. If the handler cannot correct the error, it should log the error information and return to the operating system. If the error affects only process data, it may be possible to terminate only the affected process or virtual machine. If the error affects processor state, continued use of that processor should not occur. See individual error descriptions for further guidance.
 - If [MSR0000_017A](#)[RIPV] is set, the interrupted program can be restarted reliably at the instruction pointer address pushed onto the exception handler stack if any uncorrectable error has been corrected by software. If RIPV is clear, the interrupted program cannot be restarted reliably, although it may be possible to restart it for debugging purposes. As long as PCC is clear, it may be possible to terminate only the affected process or virtual machine.
 - When logging errors, particularly those that are not recoverable, check [MSR0000_017A](#)[EIPV] to see if the instruction pointer address pushed onto the exception handler stack is related to the machine check. If EIPV is clear, the address is not ensured to be related to the error.
 - See [2.15.1.6.1 \[MCA Differentiation Between System-Fatal and Process-Fatal Errors\]](#) for more explanation on the relationship between PCC, RIPV, and EIPV.
- Exit:
 - When an exception handler is able to successfully log an error condition, clear the MC_i_STATUS registers prior to exiting the machine check handler.
 - Prior to exiting the machine check handler, be sure to clear [MSR0000_017A](#)[MCIP]. MCIP indicates that a machine check exception is in progress. If this bit is set when another machine check exception occurs in the same core, the processor enters the shutdown state.

Additional machine check handler portability can be added by having the handler use the CPUID instruction to identify the processor and its capabilities. Implementation specific software can be added to the machine check exception handler based on the processor information reported by CPUID.

In cases where sync flood is the recommended response to a particular error, a machine check exception cannot be used in lieu of the sync flood to stop the propagation of potentially bad data.

2.15.1.6.1 MCA Differentiation Between System-Fatal and Process-Fatal Errors

The bits `MCi_STATUS[PCC]`, `MSR0000_017A[RIPV]`, and `MSR0000_017A[EIPV]` form a hierarchy, used by software to determine the degree of corruption and recoverability in the system. [Table 64](#) shows how these bits are interpreted.

Table 64: Error Scope Hierarchy

PCC	RIPV	EIPV	Comments
1	0	0	Error has corrupted system state (PCC=1), and the process program cannot be restarted (RIPV=0). The error is fatal to the system.
0	1	0/1	Error is confined to the process (PCC=0), and the process program can be restarted (RIPV=1) if any uncorrectable error is corrected by software.
	0	0/1	Corruption is confined to the process (PCC=0), but the process program cannot be restarted (RIPV=0). Continued operation of this process may not be possible without intervention, however system processing or other processes can continue with appropriate software clean up.

2.15.1.7 Error Thresholding

For some types of errors, the hardware maintains counts of the number of errors. When the counter reaches a programmable threshold, an event may optionally be triggered to inform software. This is known as error thresholding. The primary purpose of error thresholding is to help software recognize an excessive rate of errors, which may indicate marginal or failing hardware. This information can be used to make decisions about deconfiguring hardware or scheduling service actions. Counts are incremented for correctable, and uncorrectable errors.

The error thresholding hardware counts only the number of errors; it is up to software to track the errors reported over time in order to determine the rate of errors. Thresholding gives error counts on groups of resources. In order to make decisions on individual resources, a finer granularity of error information, such as MCA information for specific errors, must be utilized in order to obtain more accurate counts and to limit the scope of actions to affected hardware.

Thresholding is performed for “Error Threshold Groups” identified in the list below. For all error threshold groups, some number of correctable errors is expected and normal. There are numerous factors influencing error rates, including temperature, voltage, operating speed, and geographic location. In order to accommodate the various factors, including software latency to respond and track the error thresholding, additional guard-band above the normal rates is recommended before error rates are considered abnormal for purposes of hardware action.

The {MC0, MC1, MC2, MC5} error thresholding banks maintains counters, but do not provide interrupts when the threshold is reached; these counters must be polled.

Error thresholding groups:

- LS (MC0)
 - LS errors are counted and polled via `MSR0000_0403`.
 - LS errors are listed in [Table 213 \[LS Error Signatures\]](#).
- IF (MC1)

- IF errors are counted and polled via [MSR0000_0407](#).
- IF errors are listed in [Table 216 \[IF Error Signatures\]](#).
- CU (MC2)
 - CU errors are counted and polled via [MSR0000_040B](#).
 - CU errors are listed in [Table 220 \[CU Error Signatures\]](#).
- EX (MC5)
 - EX errors are counted and polled via [MSR0000_0417](#).
 - EX errors are listed in [Table 231 \[EX Error Signatures\]](#).
- Links (MC4)
 - Link errors are counted and reported via [MSRC000_0408](#).
 - Link errors are the errors listed in [Table 222 \[NB Error Descriptions\]](#) as “L” (Cache) in the ETG (Error Threshold Group) column.
 - For a link exhibiting excessive errors, it may be possible to reduce errors by lowering the link frequency or reducing the link width (if a bad lane can be avoided). See [2.11 \[Root Complex\]](#) for details and restrictions on configuring links.

In rare circumstances, such as two simultaneous errors in the same error thresholding group, it is possible for one error not to increment the counter. In these conditions, `MCi_STATUS[Overflow]` may indicate that an overflow occurred, but the error counter may only indicate one error.

2.15.1.8 Error Diagnosis

This section describes generalized information and algorithms for diagnosing errors. The primary goal of diagnosis is to identify the failing component for repair purposes. The secondary goal is to identify the smallest possible sub-component for de-allocation, de-configuration, or design/manufacturing root cause analysis.

Indictment means identifying the part in error. The simplest form of indictment is *self-indictment*, where the bank reporting the error is also the unit in error. The next simplest form of indictment is *eyewitness indictment*, where the part in error is not the bank reporting the error, but is identified unambiguously. Both of these forms can be considered direct indictment; the information for indictment is contained in the MCA error information. If an error is not directly indicted, then identifying the part in error is more difficult and may not be an explicit part of the error log.

In general, an address logged in the MCA is useful for direct indictment only if the address identifies a physical location in error, such as a cache index. Logical addresses, while identifying the data, do not identify the location of the data.

If possible, physical storage locations in caches should be checked to determine whether the error is a soft error (a temporary upset of the stored value) or a hard fault (malfunctioning hardware). A location which has had a soft error can be corrected by writing a new value to the location; a reread of the location should see the new value. Hard faults cannot be corrected by writing a new value; the hardware persistently returns the previous value. If such checking is not possible, a grossly simplifying assumption can be made that uncorrected errors are hard and corrected errors are soft. Repeated corrected errors from the same location are an indication that the fault is actually hard.

Determining whether corrected errors represent a hard fault or a soft error requires understanding the access patterns and any attempts to correct the faulty data in place. An attempt to correct the data in place creates two *epochs*, one before the correction event and one after. If an error is seen at the same location in two different epochs (especially back-to-back epochs), it is more likely that the cause is a hard fault, since the error has persisted or repeated through an in place correction. The more epochs in which a error is seen, the higher the likelihood of it being caused by a hard fault.

As an example, consider a correctable error found during a read from DRAM. If the DRAM redirect scrubber is enabled (`D18F3x5C[ScrubReDirEn]`), the data in error is corrected in place, and this event conceptually creates a new epoch. If the original fault was due to a soft error, a read of the same data in the new epoch should not encounter a data error. If the original fault was due to a hard fault (e.g., a stuck bit), a read of the data in the new epoch will likely result in another corrected or uncorrected error.

There are numerous correction events that can be used to separate time periods into epochs. These include DRAM redirect scrubs, DRAM sequential scrubs, cache scrubs, cache writes, cache flushes, resets, and others.

2.15.1.8.1 Common Diagnosis Information

A common set of diagnosis information is useful for many problems. [Table 65](#) indicates the minimum set of generally useful diagnostic information that should be collected by software, unless the specifics of the problem are known to be narrower, based on the error code or other information.

It is useful to collect configuration information to ensure that the behavior is not caused by misconfiguration.

Table 65: Registers Commonly Used for Diagnosis

MCA Bank	Status	Configuration
MC0	MSR0000_0401 MSR0000_0402 MSR0000_0403	MSR0000_0400 MSRC001_1022 MSRC001_0044
MC1	MSR0000_0405 MSR0000_0406 MSR0000_0407	MSR0000_0404 MSRC001_0045
MC2	MSR0000_0409 MSR0000_040A MSR0000_040B	MSR0000_0408 MSRC001_0046 MSRC001_1023
MC3	Reserved	Reserved
MC4	MSR0000_0411 MSR0000_0412 MSR0000_0413 MSRC000_0408 D18F3x54 D18F2xAC	MSR0000_0410 MSRC001_0048 D18F3x40 D18F3x44 D18F3xE4 D18F3xE8 MSRC001_001F D18F3x180
MC5	MSR0000_0415 MSR0000_0416 MSR0000_0417	MSR0000_0414 MSRC001_0049
MC6	MSR0000_0419 MSR0000_041A MSR0000_041B	MSR0000_0418 MSRC001_004A

If examining MCA registers after startup, determine the cause of the startup:

- INIT; `D18F0x6C[InitDet]`.
- Cold reset; `D18F0x6C[ColdRstDet]`.

- Warm reset; if not INIT or cold reset.

To see if a link failure occurred, examine [D18F0x84\[LinkFail\]](#). If set, look for additional information:

- Receipt of a sync, such as during a sync flood, saves a status of Sync Error in MC4_STATUS.
- CRC error saves a status of CRC Error in MC4_STATUS. See [D18F0x84\[CrcErr and CrcFloodEn\]](#).
- Link not present does not save status in MC4_STATUS. See [D18F0x84\[InitComplete\]](#).

Other registers may be needed depending on the specific error symptoms.

2.15.2 Error Injection and Simulation

Error injection allows the introduction of errors into the system for test and debug purposes. See the following sections for error injection details:

- Link:
 - [D18F3x44\[GenLinkSel, GenSubLinkSel, GenCrcErrByte1, GenCrcErrByte0\]](#).

Error simulation involves creating the appearance to software that an error occurred. This is performed by manually setting the MCA registers with desired values (see [MSRC001_0015\[McStatusWrEn\]](#)), and then driving the software via INT18. McStatusWrEn can be used to debug machine check interrupt handlers. When McStatusWrEn is set, privileged software can write non-zero values to the specified registers without generating exceptions, and then simulate a machine check using the INT18 instruction (INT n instruction with an operand of 18). Setting a reserved bit in these registers does not generate an exception when this mode is enabled. However, setting a reserved bit may result in undefined behavior.

3 Registers

This section provides detailed field definitions for the core register sets in the processor.

3.1 Register Descriptions and Mnemonics

Each register in this document is referenced with a mnemonic. Each mnemonic is a concatenation of the register-space indicator and the offset of the register. Here are the mnemonics for the various register spaces:

- **IOXXX**: x86-defined input and output address space registers; XXX specifies the hexadecimal byte address of the IO instruction. This space includes IO-space configuration access registers [IOCF8 \[IO-Space Configuration Address\]](#) and [IOCF8 \[IO-Space Configuration Data Port\]](#).
- **APICXX0**: APIC memory-mapped registers; XX0 is the hexadecimal byte address offset from the base address. See [2.4.8.1.2 \[APIC Register Space\]](#).
- **CPUID FnXXXX_XXXX_EiX[_xYYYY]**: processor capabilities information returned by the CPUID instruction. See [3.17 \[CPUID Instruction Registers\]](#). Each core may only access this information for itself.
- **MSRXXXX_XXXX**: **MSRs**; XXXX_XXXX is the hexadecimal MSR number. This space is accessed through x86-defined RDMSR and WRMSR instructions. Unless otherwise specified there is one set of these registers **Per-core**. See [2.4.1 \[Compute Unit\]](#).
- **DXFYxZZZ**: PCI-defined configuration space; X specifies the hexadecimal device number (this may be 1 or 2 digits), Y specifies the function number, and ZZZ specifies the hexadecimal byte address (this may be 2 or 3 digits); e.g., D18F3x40 specifies the register at device 18h, function 3, and address 40h. See [2.7 \[Configuration Space\]](#), for details about configuration space.
 - Some register in D18F2xXXX have the `_dct[1:0]` and/or `_mp[1:0]` mnemonic suffix. See [2.9.1 \[DCT Configuration Registers\]](#).
- **IOMMUxX_XXXX**: IOMMU memory mapped registers; X_XXXX specifies the hexadecimal byte address offset (this may be 2 to 5 digits) from the base address register; The base address for this space is specified by [D0F2x44 \[IOMMU Base Address Low\]](#) and [D0F2x48 \[IOMMU Base Address High\]](#). See [3.15 \[IOMMU Memory Mapped Registers\]](#).
- **PMCxXXX**: performance monitor events; XXX is the hexadecimal event counter number programmed into [MSRC001_020\[A,8,6,4,2,0\] \[Performance Event Select \(PERF_CTL\[5:0\]\)\]\[EventSelect\]](#); See [2.6.1 \[Core Performance Monitor Counters\]](#). **NBPMCxXXX**: NB performance monitor events; XXX is the hexadecimal event counter number programmed into [MSRC001_024\[6,4,2,0\] \[Northbridge Performance Event Select \(NB_PERF_CTL\[3:0\]\)\]\[EventSelect\]](#); See [2.6.2 \[NB Performance Monitor Counters\]](#).
- When PMCxXXX or NBPMCxXXX is followed by [z:y] then UnitMask[z:y] is being specified.

Each mnemonic may specify the location of one or more registers that share the same base definition. A mnemonic that specifies more than one register will contain one or more ranges within braces. The ranges are specified as follows:

- Comma separated lists [A,B]: Define specific instances of a register, e.g., D0F3x[1,0]40 defines two registers D0F3x40 and D0F3x140.
- Colon separated ranges [A:B]: Defines all registers that contain the range between A and B. Examples:
 - D0F3x[50:40] defines five registers D0F3x40, D0F3x44, D0F3x48, D0F3x4C, and D0F3x50.
 - D[8:2]F0x40 defines seven registers D2F0x40, D3F0x40, D4F0x40, D5F0x40, D6F0x40, D7F0x40, and D8F0x40.
 - D0F0xE4_x013[2:0]_0000 defines three registers D0F0xE4_x0130_0000, D0F0xE4_x0131_0000, and D0F0xE4_x0132_0000.
- Colon separated ranges with an explicit step [A:BstepC]: Defines the registers from A to B, C defines the offset between registers., e.g., D0F3x[50:40:step8] defines three registers D0F3x40, D0F3x48, and D0F3x50.

The processor includes a single set of IO-space and configuration-space registers. However, APIC, CPUID, and MSR register spaces are implemented once per processor core. Access to IO-space and configuration space registers may require software-level techniques to ensure that no more than one core attempts to access a register at a time.

The following is terminology found in the register descriptions.

Table 66: Terminology in Register Descriptions

Term	Definition
BIOS	Software recommendation syntax. See 3.1.2 [Software Recommendation (BIOS, SBIOS, CBIOS, etc.)].
SBIOS	
CBIOS	
IBIOS	
VBIOS	
OS	
Driver	
See	Reference to remote definition.
Alias	The alias keyword allows the definition of a soft link between two registers. <ul style="list-style-type: none"> • X is an alias of Y: X is a soft link to the register Y. • X1, X2 are an alias of Y: Both X1 and X2 are soft links to Y.
IF	Allows conditional definition as a function of register fields. The syntax is: <ul style="list-style-type: none"> • IF (conditional-expression) THEN definition ENDIF. • IF (conditional-expression) THEN definition ELSE definition ENDIF. • IF (conditional-expression) THEN definition ELSEIF (conditional-expression) definition ELSE definition ENDIF.
THEN	
ELSEIF	
ELSE	
ENDIF	
Access Types	
Read	Capable of being read by software.
Read-only	Capable of being read but not written by software.
Write	Capable of being written by software.
Write-only	Write-only. Capable of being written by software. Reads are undefined.
Read-write	Capable of being written by software and read by software.
Set-by-hardware	Register field is set high by hardware, set low by hardware, or updated by hardware.
Cleared-by-hardware	
Updated-by-hardware	
Write-1-to-clear	Software must write a 1 to the bit in order to clear it. Writing a 0 to these bits has no affect.
Write-1-only	Software can set the bit high by writing a 1 to it. Writes of 0 have no effect.
SharedC	Shared by both cores of a compute unit: SharedC (shared coherent) or SharedNC (shared non-coherent). See 2.4.1.1 [Registers Shared by Cores in a Compute Unit] for a definition of coherent.
SharedNC	
Reset-applied	Takes affect on warm reset.
GP-read	GP exception occurs on read.

Table 66: Terminology in Register Descriptions

Term	Definition
GP-write	GP exception occurs on write.
GP-read-write	GP exception occurs on a read or a write.
Per-core	One instance per core. Only valid for MMIO config space. Writes of these bits from one core only affect that core's register. Reads return the values appropriate to that core.
Per-compute-unit	One instance per compute unit. Writes of these bits from one core only affect that compute unit's register. Reads return the values appropriate to that compute unit.
Per-node	One instance per node. See 3.1.1 [Northbridge MSRs In Multi-Core Products].
Not-same-for-all	Provide indication as to whether all instances of a given register should be the same across all cores/nodes according to the following equation:
Same-for-all	SameOnAllCheckEnabled = (Writable && (same-for-all MSR) && ~(not-same-for-all UpdatedByHw)). UpdatedByHw = (Updated-by-hardware set-by-hardware cleared-by-hardware set-when-done cleared-when-done).
Field Definitions	
Reserved	Field is reserved for future use. Software is required to preserve the state read from these bits when writing to the register. Software may not depend on the state of reserved fields nor on the ability of such fields to return the state previously written.
Unused	Field is reserved for future use. Software is not required to preserve the state read from these bits when writing to the register. Software may not depend on the state of unused fields nor on the ability of such fields to return the state previously written.
MBZ	Must be zero. If software attempts to set an MBZ bit to 1, a general-protection exception (#GP) occurs.
RAZ	Read as zero. Writes are ignored, unless RAZ is combined with write, write-1-only or write-once.
Reset Definitions	
Reset	The reset value of each register is provided below the mnemonic or in the field description. Unless otherwise noted, the register state matches the reset value when RESET_L is asserted (either a cold or a warm reset). Reset values may include: <ul style="list-style-type: none"> • X: an X in the reset value indicates that the field resets (warm or cold) to an unspecified state.
Cold reset	The field state is not affected by a warm reset (even if the field is labeled "cold reset: X"); it is placed into the reset state when PWROK is deasserted. See "Reset" above for the definition of characters that may be found in the cold reset value.
Value	The current value of a read-only field or register. A value statement explicitly defines the field or register as read-only and the value returned under all conditions including after reset events. A field labeled "Value:" will not have a separate reset definition.

3.1.1 Northbridge MSRs In Multi-Core Products

MSRs that control Northbridge functions are shared between all cores on the node in a multi-core processor (e.g. MSR0000_0410). If control of Northbridge functions is shared between software on all cores, software must ensure that only one core at a time is allowed to access the shared MSR. Some MSR's are conditionally shared; see D18F3x44[NbMcaToMstCpuEn].

3.1.2 Software Recommendation (BIOS, SBIOS, CBIOS, etc.)

The following keywords specify the recommended value to be set by software.

- BIOS: AMD BIOS.
- SBIOS: Platform BIOS.
- CBIOS: Chip-set BIOS.
- OS: Operating system.
- Driver: Device driver software settings invoked by the OS.

Syntax: BIOS:<integer-expression>. Any of the supported tags can be substituted for BIOS.

If “BIOS:” occurs in a register field then the recommended value is applied to the field. If “BIOS:” occurs after a register name but outside of a register field table row then the recommended value is applied to the width of the register.

3.1.3 Mapping Tables

The following mapping table types are defined.

3.1.3.1 Register Mapping

The register mapping table specifies the specific function for each register in a range of registers.

[Table 182](#), for example, specifies that the D18F5x160 function is for NB P-state 0.

3.1.3.2 Index Mapping

The index mapping table is similar to the register mapping table, but specifies the register by index instead of by full register mnemonic.

[Table 146](#), for example, specifies that the D18F2x98_dct[1:0][31:0]==0D0F_0002h, or D18F2x9C_x0D0F_0002_dct[1:0], function is for Byte 0.

3.1.3.3 Field Mapping

The field mapping table maps the fields of a range of registers. The rows are the registers that are mapped. Each column specifies a field bit range that is mapped by that column for all registers. The cell at the intersection of the register and the field bit range specifies the suffix that is appended to the register field. “Reserved” specifies that the field is reserved for the register of that row.

[Table 136](#), for example, specifies that the fields at D18F2x9C_x0000_0[3:0]01[31:24] should have the suffix Byte3, resulting in WrDatGrossDlyByte3 and WrDatFineDlyByte3.

3.1.3.4 Broadcast Mapping

The broadcast mapping table maps a register address to a range of register addresses. The register address is formed by the concatenation of the row address with the column address. The cell at the intersection of the row and column address is a range of register addresses that will be read or written as a group when the row and column address is read or written.

[Table 163](#), for example, specifies that a read or write to D18F2x98_dct[1:0][31:0]==0D0F_0F31h will result

in a broadcast read or write to the D18F2x9C_x0D0F_0[8:0]31 range of registers.

3.1.3.5 Valid Values

The valid values table defines the valid values for one or more register fields. The valid values table is equivalent in function to the Bits/Description tables in register fields (E.g. D18F0x16C[ForceFullT0]) and is most often used when the table becomes too large and unwieldy to be included into the register field. (E.g. Table 132 [Memory Clock Frequency Value Definition])

3.2 IO Space Registers

See 3.1 [Register Descriptions and Mnemonics] for a description of the register naming convention.

IOCF8 IO-Space Configuration Address

Reset: 0. **IOCF8** [IO-Space Configuration Address], and **IOCFC** [IO-Space Configuration Data Port], are used to access system configuration space, as defined by the PCI specification. **IOCF8** provides the address register and **IOCFC** provides the data port. Software sets up the configuration address by writing to **IOCF8**. Then, when an access is made to **IOCFC**, the processor generates the corresponding configuration access to the address specified in **IOCF8**. See 2.7 [Configuration Space].

IOCF8 may only be accessed through aligned, DW IO reads and writes; otherwise, the accesses are passed to the appropriate IO link. Accesses to **IOCF8** and **IOCFC** received from an IO link are treated as all other IO transactions received from an IO link and are forwarded based on the settings in D18F1x[DC:C0] [IO-Space Base/Limit]. **IOCF8** and **IOCFC** in the processor are not accessible from an IO link.

Bits	Description
31	ConfigEn: configuration space enable. Read-write. 1=IO read and write accesses to IOCFC are translated into configuration cycles at the configuration address specified by this register. 0=IO read and write accesses are passed to the appropriate IO link and no configuration access is generated.
30:28	Reserved.
27:24	ExtRegNo: extended register number. Read-write. ExtRegNo provides bits[11:8] and RegNo provides bits[7:2] of the byte address of the configuration register. ExtRegNo is reserved unless it is enabled by MSRC001_001F [EnableCf8ExtCfg].
23:16	BusNo: bus number. Read-write. Specifies the bus number of the configuration cycle.
15:11	Device: device number. Read-write. Specifies the device number of the configuration cycle.
10:8	Function. Read-write. Specifies the function number of the configuration cycle.
7:2	RegNo: register address. Read-write. See IOCF8 [ExtRegNo].
1:0	Reserved.

IOCFC IO-Space Configuration Data Port

Bits	Description
31:0	Data. Read-write. Reset: 0. See IOCF8 .

3.3 Device 0 Function 0 (Root Complex) Configuration Registers

See 3.1 [Register Descriptions and Mnemonics] for a description of the register naming convention. See 2.7 [Configuration Space] for details about how to access this space.

D0F0x00 Device/Vendor ID

Reset: 1410_1022h.

Bits	Description
31:16	DeviceID: device ID. Read-only.
15:0	VendorID: vendor ID. Read-only.

D0F0x04 Status/Command

Reset: 0220_0004h.

Bits	Description
31	ParityErrorDetected: parity error detected. Read-only.
30	SignaledSystemError: signaled system error. Read; write-1-to-clear. 1=FCH generated a system error.
29	ReceivedMasterAbort: received master abort. Read; write-1-to-clear.
28	ReceivedTargetAbort: received target abort. Read; write-1-to-clear.
27	SignalTargetAbort: signaled target abort. Read-only.
26:25	DevselTiming: DEVSEL# Timing. Read-only.
24	Reserved.
23	FastBackCapable: fast back-to-back capable. Read-only.
22	Reserved.
21	PCI66En: 66 MHz capable. Read-only.
20	CapList: capability list. Read-only. 1=Capability list supported.
19:10	Reserved.
9	FastB2BEn: fast back-to-back enable. Read-write.
8	SerrEn: system error enable. Read-write.
7	Reserved.
6	ParityErrorEn: parity error response enable. Read-only.
5	PalSnoopEn: VGA palette snoop enable. Read-only.
4	MemWriteInvalidateEn: memory write and invalidate enable. Read-only.
3	SpecialCycleEn: special cycle enable. Read-only.
2	BusMasterEn: bus master enable. Read-only.
1	MemAccessEn: memory access enable. Read-write. This bit controls if memory accesses by this device are accepted or not. 1=Enabled. 0=Disabled.
0	IoAccessEn: IO access enable. Read-only.

D0F0x08 Class Code/Revision ID

Reset: 0600_0000h.

Bits	Description
31:8	ClassCode: class code. Read-only. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only.

D0F0x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:24	BIST. Read-only.
23:16	HeaderTypeReg. Read-only. 00h=Single function device.
15:8	LatencyTimer. Read-write.
7:0	CacheLineSize. Read-only.

D0F0x2C Subsystem and Subvendor ID

Reset: 1410_1022h.

Bits	Description
31:16	SubsystemID. Read-only.
15:0	SubsystemVendorID. Read-only.

D0F0x34 Capabilities Pointer

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7:0	CapPtr: capabilities pointer. Read-only. There is no capability list.

D0F0x4C PCI Control

Reset: 0000_1002h.

Bits	Description
31:27	Reserved.
26	HPDis: hot plug message disable. Read-write. 1=Hot plug message generation is disabled.
25:24	Reserved.
23	MMIOEnable: memory mapped IO enable. Read-write. 1=Decoding of MMIO cycles is enabled.
22:15	Reserved.

14:12	CfgRdTime . Read-write. Specifies the propagation delay for read data on the configuration bus. <u>Bits</u> <u>Definition</u> 111b-000b <7-CfgRdTime> Clocks.
11	CRS: configuration request retry detected . Read; Write-1-to-clear. 1=Configuration request retry was detected.
10:6	Reserved.
5	SerrDis: system error message disable . Read-write. 1=The generation of SERR messages is disabled.
4	PMEDis: PME disable . Read-write. 1=The generation of PME messages is disabled.
3	Cf8Dis: Cf8 disable . Read-write. 1=Configuration accesses through IO address Cf8h to this device are disabled.
2	Reserved.
1	ApicEnable: APIC enable . Read-write. 1=APIC is enabled.
0	Function1Enable: device 0 function 1 enable . Read-write. 1=Configuration accesses to device 0 function 1 are enabled.

D0F0x60 Miscellaneous Index

Reset: 0000_0000h.

The index/data pair registers [D0F0x60](#) and [D0F0x64](#) is used to access the registers [D0F0x64_x\[FF:00\]](#). To read or write to one of these register, the address is written first into the address register [D0F0x60](#) and then the data are read or written by read or write the data register [D0F0x64](#).

Bits	Description
31:8	Reserved.
7	MiscIndWrEn: miscellaneous index write enable . Read-write. If set writes to D0F0x64 are enabled.
6:0	MiscIndAddr: miscellaneous index register address . Read-write.

D0F0x64 Miscellaneous Index Data

See [D0F0x60](#).

D0F0x64_x00 Northbridge Control

Reset: 0000_0002h.

Bits	Description
31:8	Reserved.
7	HwInitWrLock . Read-write. 1=Lock HWInit registers. 0=Unlock HWInit registers.
6	NbFchCfgEn: device 8 enable . Read-write. 1=Bridge device 8 is enabled.
5:0	Reserved.

D0F0x64_x0B IOC Link Control

Reset: 0000_0000h.

Bits	Description
31:24	Reserved.
23	IocFchSetPmeTurnOffEn. Read-write. 1=Enables the PME_Turn_Off/PME_To_Ack mechanism between northbridge and FCH.
22	Reserved.
21	IocFchSetPowEn: set slot power enable. Read-write. 1=Enables sending set_slot_power_limit/scale messages to the FCH.
20	SetPowEn: set slot power enable. Read-write. 1=Enables sending set_slot_power messages to the FCH.
19:0	Reserved.

D0F0x64_x0C IOC Bridge Control

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7	Dev7BridgeDis. Read-write. 1=Bus 0, device 7 bridge functionality is hidden.
6	Dev6BridgeDis. Read-write. 1=Bus 0, device 6 bridge functionality is hidden.
5	Dev5BridgeDis. Read-write. 1=Bus 0, device 5 bridge functionality is hidden.
4	Dev4BridgeDis. Read-write. 1=Bus 0, device 4 bridge functionality is hidden.
3	Dev3BridgeDis. Read-write. 1=Bus 0, device 3 bridge functionality is hidden.
2	Dev2BridgeDis. Read-write. 1=Bus 0, device 2 bridge functionality is hidden.
1:0	Reserved.

D0F0x64_x0D IOC PCI Configuration

Reset: 0000_0001h.

Bits	Description
31:1	Reserved.
0	PciDev0Fn2RegEn. Read-write. BIOS: See 2.12.2 [IOMMU Initialization] . 1=Enable configuration accesses to device 0 function 2.

D0F0x64_x16 IOC Advanced Error Reporting Control

Reset: 0000_0001h.

Bits	Description
31:1	Reserved.
0	AerUrMsgEn: AER unsupported request message enable. Read-write. 1=AER unsupported request messages are enabled.

D0F0x64_x19 Top of Memory 2 Low

Reset: 0000_0000h.

Bits	Description
31:23	Tom2[31:23]: top of memory 2. Read-write. BIOS: MSRC001_001D [Tom2[31:23]]. This field specifies the maximum system address for upstream read and write transactions that are forwarded to the host bridge. All addresses less than or equal to this system address are forwarded to DRAM and are not checked to determine if the transaction is a peer-to-peer transaction. All upstream reads with addresses greater than this system address are master aborted.
22:1	Reserved.
0	TomEn: top of memory enable. Read-write. BIOS: MSRC001_0010 [MtrrTom2En]. 1=Top of memory check enabled.

D0F0x64_x1A Top of Memory 2 High

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7:0	Tom2[39:32]: top of memory 2. Read-write. BIOS: MSRC001_001D [Tom2[39:32]]. See D0F0x64_x19 [Tom2].

D0F0x64_x1C Internal Graphics PCI Control 1

Reset: 0080_0000h.

Bits	Description
31:24	Reserved.
23	RcieEn. IF (D0F0x64_x1C [WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 1. 1=Root complex integrated endpoint mode. 0=Legacy PCI device mode.
22:18	Reserved.
17	F0En. IF (D0F0x64_x1C [WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 1. 1=Internal graphics enabled. 0=Internal Graphics disabled.
16	IoBarDis. IF (D0F0x64_x1C [WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. 1=Graphics IO base address register disabled. 0= Graphics IO base address register enabled.
15:12	Reserved.

11	Audio64BarEn. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. Controls the size of the audio BAR. 1=64-bit BAR. 0=32-bit BAR.																				
10	AudioNonlegacyDeviceTypeEn. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 0. 1=PCIe device. 0=Legacy PCI device.																				
9	MsiDis. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 0. 1=MSI interrupts disabled. 0=MSI interrupts enabled.																				
8	AudioEn. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 1. 1=HDMI™ audio enabled. 0=HDMI audio disabled.																				
7	Reserved.																				
6	RegApSize. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 1. Specifies the size of the graphics register aperture. 0=64KB. 1=256KB.																				
5:3	MemApSize. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. Specifies the size of the frame buffer aperture. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>128MB</td> <td>100b</td> <td>512MB</td> </tr> <tr> <td>001b</td> <td>256MB</td> <td>101b</td> <td>1GB</td> </tr> <tr> <td>010b</td> <td>64MB</td> <td>110b</td> <td>2GB</td> </tr> <tr> <td>011b</td> <td>32MB</td> <td>111b</td> <td>4GB</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	000b	128MB	100b	512MB	001b	256MB	101b	1GB	010b	64MB	110b	2GB	011b	32MB	111b	4GB
Bits	Definition	Bits	Definition																		
000b	128MB	100b	512MB																		
001b	256MB	101b	1GB																		
010b	64MB	110b	2GB																		
011b	32MB	111b	4GB																		
2	F064BarEn. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. 1=64-bit base address registers. 0=32-bit base address registers.																				
1	F0NonlegacyDeviceTypeEn. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 0. 1=PCIe device. 0=Legacy PCI device.																				
0	WriteDis. IF (D0F0x64_x1C[WriteDis]==1) THEN Read-only. ELSE Read-write. ENDIF. 1=Register is read-only and the GPU PCI interface is configured. 0=Register is read-write.																				

D0F0x64_x1D Internal Graphics PCI Control 2

Reset: 0000_0001h.

Bits	Description
31:4	Reserved.
3	Vga16En: VGA IO 16 bit decoding enable. Read-write. 1=Address bits 15:10 for VGA IO cycles are decoded. 0=Address bits 15:10 for VGA IO cycles are ignored.
2	Reserved.
1	VgaEn: VGA enable. Read-write. Affects the response by the internal graphics to compatible VGA addresses when IntGfxAsPcieEn=1. 1=The internal graphics decodes the following accesses: <ul style="list-style-type: none"> • Memory accesses in the range of A0000h to BFFFFh. • IO address where address bits 9:0 are in the ranges of 3B0h to 3BBh or 3C0h to 3DFh. For IO cycles the decoding of address bits 15:10 depends on Vga16En.
0	IntGfxAsPcieEn: internal graphics is RC integrated device. Read-write. BIOS: 1. 1=Integrated graphics is device 1 on bus 0 and operates as a RC integrated device. 0=Integrated graphics is located behind a PCI-to-PCI bridge and is device 5 on the bus behind the bridge. The bridge device is device 1 on bus 0.

D0F0x64_x20 Programmable Device Remap 0

Reset: BA97_6542h.

Bits	Description
31:20	Reserved.
19:16	GppPortEDevmap . Read-write. Maps any downstream PCIE bridge, except SB, to port E of PCIE GPPSB core.
15:12	GppPortDDevmap . Read-write. Maps any downstream PCIE bridge, except SB, to port D of PCIE GPPSB core.
11:8	GppPortCDevmap . Read-write. Maps any downstream PCIE bridge, except SB, to port C of PCIE GPPSB core.
7:4	GppPortBDevmap . Read-write. Maps any downstream PCIE bridge, except SB, to port B of PCIE GPPSB core.
3:2	Reserved.
1	IocPcieDevRemapDis . Read-write. 1=Disables remapping.
0	ProgDevMapEn . Read-write. 1=Enable software device remapping.

D0F0x64_x21 Programmable Device Remap 1

Reset: 0003_2EDCh.

Bits	Description
31:20	Reserved.
19:16	GfxPortBDevmap . Read-write. Maps any downstream PCIE bridge, except SB, to port B of PCIE GFX core.
15:12	GfxPortADevmap . Read-write. Maps any downstream PCIE bridge, except SB, to port A of PCIE GFX core.
11:0	Reserved.

D0F0x64_x22 LCLK Control 0

Reset: 7F3F_8100h.

Bits	Description
31	Reserved.
30	SoftOverrideClk0 . Read-write. BIOS: 0. 1=Dynamic clock gating disabled for the host request path to the PCIe cores.
29	SoftOverrideClk1 . Read-write. BIOS: 0. 1=Dynamic clock gating disabled for the host request path to the internal graphics and the host response path.
28	SoftOverrideClk2 . Read-write. BIOS: 0. 1=Dynamic clock gating disabled for the host configuration requests.
27	SoftOverrideClk3 . Read-write. BIOS: 0. 1=Dynamic clock gating disabled for the debug bus path.

26	SoftOverrideClk4. Read-write. BIOS: 0. 1=Dynamic clock gating disabled for the host request path to the configuration block.
25:0	Reserved.

D0F0x64_x23 LCLK Control 1

Reset: 7F3F_8100h.

Bits	Description
31	Reserved.
30	SoftOverrideClk0. Read-write. BIOS: IF (D0F2x44[IommuEnable]==1) THEN 1 ELSE 0 ENDIF. 1=Dynamic clock gating disabled for upstream DMA requests from all sources.
29	SoftOverrideClk1. Read-write. BIOS: IF (D0F2x44[IommuEnable]==1) THEN 1 ELSE 0 ENDIF. 1=Dynamic clock gating disabled for upstream DMA requests from the GPPFCH link core.
28	SoftOverrideClk2. Read-write. BIOS: 0. 1=Dynamic clock gating disabled for upstream DMA requests from internal graphics and its DMA response reordering path.
27	SoftOverrideClk3. Read-write. BIOS: IF (D0F2x44[IommuEnable]==1) THEN 1 ELSE 0 ENDIF. 1=Dynamic clock gating disabled for upstream DMA requests from internal graphics.
26	SoftOverrideClk4. Read-write. BIOS: IF (D0F2x44[IommuEnable]==1) THEN 1 ELSE 0 ENDIF. 1=Dynamic clock gating disabled for upstream DMA requests from the Gfx link core.
25:0	Reserved.

D0F0x64_x46 IOC Features Control

Reset: 0001_3063h.

Bits	Description						
31:17	Reserved.						
16	Msi64bitEn: 64 bit MSI enable. Read-write. 1=64-bit MSI support enabled. 0=64-bit MSI support disabled.						
15:3	Reserved.						
2:1	P2PMode: peer-to-peer mode. Read-write. BIOS: 00b. Specifies how upstream write transactions above D0F0x64_x19[Tom2] are completed. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Master abort writes that do not hit one of the internal PCI bridges. Forward writes that hit one of the internal PCI bridges to the bridge.</td> </tr> <tr> <td>11b-01b</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Definition	00b	Master abort writes that do not hit one of the internal PCI bridges. Forward writes that hit one of the internal PCI bridges to the bridge.	11b-01b	Reserved.
Bits	Definition						
00b	Master abort writes that do not hit one of the internal PCI bridges. Forward writes that hit one of the internal PCI bridges to the bridge.						
11b-01b	Reserved.						
0	Reserved.						

D0F0x64_x5[B,9,7,5,3,1] IOC PCIe Device Control

Reset: 0000_0000h.

Table 67: BIOS recommendations for D0F0x64_x5[B,9,7,5,3,1]

D0F0x60 [6:0]	Function	BIOS
51h	Device 2	IF (D0F0x64_x0C[Dev2BridgeDis]==0) THEN 0010_0000h ELSE 0000_0000h ENDIF.
53h	Device 3	IF (D0F0x64_x0C[Dev3BridgeDis]==0) THEN 0010_0000h ELSE 0000_0000h ENDIF.
55h	Device 4	IF (D0F0x64_x0C[Dev4BridgeDis]==0) THEN 0010_0000h ELSE 0000_0000h ENDIF.
57h	Device 5	IF (D0F0x64_x0C[Dev5BridgeDis]==0) THEN 0010_0000h ELSE 0000_0000h ENDIF.
59h	Device 6	IF (D0F0x64_x0C[Dev6BridgeDis]==0) THEN 0010_0000h ELSE 0000_0000h ENDIF.
5Bh	Device 7	IF (D0F0x64_x0C[Dev7BridgeDis]==0) THEN 0010_0000h ELSE 0000_0000h ENDIF.

Bits	Description
31:21	Reserved.
20	SetPowEn: set slot power enable. Read-write. 1=Enables the set_slot_power message to the FCH.
19:0	Reserved.

D0F0x7C IOC Configuration Control

Reset: 0000_0000h.

Bits	Description
31:1	Reserved.
0	ForceIntGFXDisable: internal graphics disable. Read-write. Setting this bit disables bridge device 1 on bus 0 and all devices behind this bridge.

D0F0x84 Link Arbitration

Reset: 0300_0018h.

Bits	Description
31:10	Reserved.
9	PmeTurnOff: PME_Turn_Off message trigger. Read-write. 1=Trigger a PME_Turn_Off message to all downstream devices if PmeMode=1.
8	PmeMode: PME message mode. Read-write. 1=PME_Turn_Off message is triggered by writing PmeTurnOff. 0=PME_Turn_Off message is triggered by a message from the FCH.
7:5	Reserved.
4	Ev6Mode: EV6 mode. Read-write. BIOS: 1. 1=The links decode the memory range from 640K to 1M.
3	VgaHole: vga memory hole. Read-write. This bit creates a hole in memory for the VGA memory range. 1=Requests hitting the VGA range are checked against PCI bridge memory ranges instead of being forwarded to system memory.
2:0	Reserved.

D0F0x90 Northbridge Top of Memory

Reset: 0000_0000h.

Bits	Description
31:23	TopOfDram. Read-write. BIOS: MSRC001_001A [TOM[31:23]]. Specifies the address that divides between MMIO and DRAM. From TopOfDram to 4G is MMIO; below TopOfDram is DRAM. See 2.4.3 [Access Type Determination] .
22:0	Reserved.

D0F0x94 Northbridge ORB Configuration Offset

Reset: 0000_0000h.

The index/data pair [D0F0x94](#) and [D0F0x98](#) are used to access [D0F0x98_x\[FF:00\]](#). To read or write to one of these register, the address is written first into the address register [D0F0x94](#) and then the data are read or written by read or write the data register [D0F0x98](#).

Bits	Description
31:9	Reserved.
8	OrbIndWrEn: ORB index write enable. Read-write. 1=Writes to D0F0x98 are enabled.
7	Reserved.
6:0	OrbIndAddr: ORB index register address. Read-write.

D0F0x98 Northbridge ORB Configuration Data PortSee [D0F0x94](#).**D0F0x98_x06 ORB Downstream Control 0**

Reset: 0000_0000h.

Bits	Description
31:27	Reserved.
26	UmiNpMemWrEn. Read-write. BIOS: See 2.11.4.1 . 1=NP protocol over UMI for memory-mapped writes targeting LPC enabled. This bit may be set to avoid a deadlock condition.
25:0	Reserved.

D0F0x98_x07 ORB Upstream Arbitration Control 0

Reset: 0000_0000h.

Bits	Description
31	SMUCsrIsocEn. Read-write. BIOS: 1. 1=CSR accesses go through ISOC channel. If this bit is set, D0F0x98_x1E [HiPriEn] must also be set.
30	UnadjustThrottlingStpclk. Read-write. BIOS: 1. 1=Enable normal UnitID for STPCLK message.
29:16	Reserved.

15	DropZeroMaskWrEn. Read-write. BIOS: 1. 1=Drop byte write request that have all bytes masked. 0=Forward byte write request that have all bytes masked.						
14	MSIHTIntConversionEn. Read-write. BIOS: 0. 1=MSI to HT interrupt conversion enabled.						
13:8	Reserved.						
7	IommuIsocPassPWMode. Read-write. BIOS: 1. 1=Always set PassPW for IOMMU upstream isochronous requests.						
6	DmaReqRespPassPWMode. Read-write. BIOS: 0. Specifies the RespPassPW bit for non-posted upstream DMA requests. <table border="1" data-bbox="267 514 803 619"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Always 1.</td> </tr> <tr> <td>1</td> <td>Value passed from IOC.</td> </tr> </tbody> </table>	Bit	Description	0	Always 1.	1	Value passed from IOC.
Bit	Description						
0	Always 1.						
1	Value passed from IOC.						
5	Reserved.						
4	IommuBwOptEn. Read-write. BIOS: 1. 1=Optimize IOMMU L2 byte write by detecting consecutive DW mask and translate the request to DW write.						
3:1	Reserved.						
0	IocBwOptEn. Read-write. BIOS: 1. 1=Enable optimization of byte writes by detecting consecutive DW masks and translating the request to DW writes.						

D0F0x98_x08 ORB Upstream Arbitration Control 1

Reset: 0008_0808h. BIOS: 0001_0808h.

This register specifies the weights of the weighted round-robin arbiter in stage 1 of the upstream arbitration for non-posted reads.

Bits	Description
31:24	Reserved.
23:16	NpWrrLenC. Read-write. This field defines the maximum number of non-posted read requests from the GPU that are serviced before the arbiter switches to the next client.
15:8	NpWrrLenB. Read-write. This field defines the maximum number of non-posted read requests from IOMMU that are serviced before the arbiter switches to the next client.
7:0	NpWrrLenA. Read-write. This field defines the maximum number of non-posted read requests from IOC that are serviced before the arbiter switches to the next client.

D0F0x98_x09 ORB Upstream Arbitration Control 2

Reset: 0000_0808h.

This register specifies the weights of the weighted round-robin arbiter in stage 1 of the upstream arbitration for posted writes.

Bits	Description
31:16	Reserved.

15:8	PWrrLenB . Read-write. This field defines the maximum number of posted write requests from the IOMMU that are serviced before the arbiter switches to the next client.
7:0	PWrrLenA . Read-write. This field defines the maximum number of posted write requests from the IOC that are serviced before the arbiter switches to the next client.

D0F0x98_x0C ORB Upstream Arbitration Control 5

Reset: 0000_0808h.

This register specifies the weights of the weighted round-robin arbiter in stage 2 of the upstream arbitration.

Bits	Description
31	Reserved.
30	StrictSelWinnerEn . Read-write. BIOS: 1. 1=Select arbitration winner when TX is idle and the FIFO is not full. 0=Select arbitration winner when TX is idle.
29:16	Reserved.
15:8	GcmWrrLenB . Read-write. BIOS: 08h. This field defines the maximum number of posted write requests from stage 1 that are getting serviced in the round-robin before the stage 2 arbiter switches to the next client.
7:0	GcmWrrLenA . Read-write. BIOS: 08h. This field defines the maximum number of non-posted read requests from stage 1 that are getting serviced in the round-robin before the stage 2 arbiter switches to the next client.

D0F0x98_x1E ORB Receive Control 0

Reset: 0000_0000h.

Bits	Description
31:2	Reserved.
1	HiPriEn . Read-write. BIOS: 1. 1=High priority channel enabled. See D0F0x98_x27[IOMMUUrAddr[31:6]] . IF (D0F0x98_x1E[HiPriEn] ==0) THEN (D0F0x98_x07[SMUCsrIsocEn] ==0) && (IOMMUx18[Isoc] ==0).
0	Reserved.

D0F0x98_x26 ORB IOMMU Control 0

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7:0	IOMMUUrAddr[39:32] . Read-write. See: D0F0x98_x27[IOMMUUrAddr[31:6]] .

D0F0x98_x27 ORB IOMMU Control 1

Reset: 0000_0000h.

Bits	Description
31:6	IOMMUUrAddr[31:6] . Read-write. BIOS: IOMMUUrAddr[39:6] must be programmed to a safe system memory address when D0F0x98_x1E[HiPriEn]=1 . IOMMUUrAddr[39:6] = { D0F0x98_x26[IOMMUUrAddr[39:32]] , IOMMUUrAddr[31:6]}. IOMMU requests that are not directed to system memory are redirected to IOMMUUrAddr.
5:0	Reserved.

D0F0x98_x28 ORB Transmit Control 0

Reset: 0000_0000h.

Bits	Description
31:2	Reserved.
1	ForceCoherentIntr . Read-write. BIOS: 1. 1=Interrupt request are forced to have coherent bit set.
0	Reserved.

D0F0x98_x2C ORB Clock Control

Reset: 000F_0000h.

Bits	Description
31:16	WakeHysteresis . Read-write. BIOS: 64h. Specifies the amount of time hardware waits after ORB becomes idle before deasserting the wake signal to the NB. Wait time = WakeHysteresis * 50ns. Values less than 64h may result in undefined behavior. Changes to this field should be done prior to setting DynWakeEn.
15:11	Reserved.
10	NBOutbWakeMask . Read-write. BIOS: 1. 0=OnOutbWake state affects OnInbWake state. 1=OnOutbWake state is masked out.
9	SBDmaActiveMask . Read-write. BIOS: 0. 0=SB_DMA_ACTIVE_L state affects OnInbWake state. 1= SB_DMA_ACTIVE_L state is masked out.
8	OrbRxIdlesMask . Read-write. BIOS: 1. 0=ORB RX idle states affect OnInbWake state. 1=ORB RX idle states are masked out.
7:2	Reserved.
1	DynWakeEn . Read-write. BIOS: 1. 1=Enable dynamic toggling of the wake signal between ORB and NB. 0=Disable dynamic toggling of the wake signal. See WakeHysteresis.
0	Reserved.

D0F0x98_x3A ORB Source Tag Translation Control 2

Reset: 0000_0000h.

Clumping allows device 2 port to have up to 64 outstanding non-posted requests when device 3 is not enabled.

Bits	Description
31:0	ClumpingEn. Read-write. Valid for clumping internal unit IDs 2 and 3 and/or unit IDs 14, 15, 16, and 17.

D0F0x98_x4[A,9] ORB LCLK Clock Control 1-0

Reset: 7F3F_8100h.

Bits	Description
31	Reserved.
30	SoftOverrideClk0. Read-write. BIOS: IF (REG==D0F0x98_x4A) THEN 1 ELSE 0 ENDIF. See: SoftOverrideClk6.
29	SoftOverrideClk1. Read-write. BIOS: 0. See: SoftOverrideClk6.
28	SoftOverrideClk2. Read-write. BIOS: 0. See: SoftOverrideClk6.
27	SoftOverrideClk3. Read-write. BIOS: 0. See: SoftOverrideClk6.
26	SoftOverrideClk4. Read-write. BIOS: 0. See: SoftOverrideClk6.
25	SoftOverrideClk5. Read-write. BIOS: 0. See: SoftOverrideClk6.
24	SoftOverrideClk6. Read-write. BIOS: 0. 1=Clock gating disabled. 0=Clock gating enabled.
23:0	Reserved.

D0F0xB8 SMU Index Address

The index/data pair [D0F0xB8](#) and [D0F0xBC](#) are used to access [D0F0xBC_x\[FFFFFFFF:00000000\]](#). To read or write to one of these register, the address is written first into the address register [D0F0xB8](#) and then the data are read or written by read or write to the data register [D0F0xBC](#).

Bits	Description
31:0	NbSmuIndAddr: smu index address. Read-write. Reset: 0.

D0F0xBC SMU Index Data

See [D0F0xB8](#).

D0F0xBC_x1F100 SCLK DPM Control

Reset: xxxx_xxxxh. See [2.5.6.1.2 \[SCLK DPM\]](#).

Bits	Description
31:25	Reserved.
24	VoltageChgEn. Read-write
23:16	SclkDpmBootState. Read-write.
15:1	Reserved. Read-write.
0	SclkDpmEn. Read-write.

D0F0xBC_x1F2[E0:00:step20] LCLK DPM Control 0

Reset: xxxx_xxxxh. Each register in [D0F0xBC_x1F2\[E0:00:step20\]](#) corresponds to one LCLK DPM state as follows.

Table 68: Register Mapping for [D0F0xBC_x1F2\[E0:00:step20\]](#)

Register	Function	Register	Function
D0F0xBC_x1F200	State 0	D0F0xBC_x1F280	State 4
D0F0xBC_x1F220	State 1	D0F0xBC_x1F2A0	State 5
D0F0xBC_x1F240	State 2	D0F0xBC_x1F2C0	State 6
D0F0xBC_x1F260	State 3	D0F0xBC_x1F2E0	State 7

Bits	Description
31:24	LowVoltageReqThreshold. Read-write.
23:16	VID. Read-write. Specifies the VDDNB VID for this DPM state.
15:8	LclkDivider. Read-write. Specifies the LCLK divisor for this DPM state.
7:1	Reserved.
0	StateValid. Read-write. 1=DPM state is valid. 0=DPM state is invalid.

D0F0xBC_x1F2[E8:08:step20] LCLK DPM Control 2

Reset: xxxx_xxxxh. Each register in [D0F0xBC_x1F2\[E8:08:step20\]](#) corresponds to one LCLK DPM state as follows.

Table 69: Register Mapping for [D0F0xBC_x1F2\[E8:08:step20\]](#)

Register	Function	Register	Function
D0F0xBC_x1F208	State 0	D0F0xBC_x1F288	State 4
D0F0xBC_x1F228	State 1	D0F0xBC_x1F2A8	State 5
D0F0xBC_x1F248	State 2	D0F0xBC_x1F2C8	State 6
D0F0xBC_x1F268	State 3	D0F0xBC_x1F2E8	State 7

Bits	Description
31:16	ResidencyCounter. Read-write.
15:8	HysteresisDown. Read-write.
7:0	HysteresisUp. Read-write.

D0F0xBC_x1F2[F0:10:step20] LCLK DPM Activity Thresholds

Reset: xxxx_xxxxh. Each register in [D0F0xBC_x1F2\[F0:10:step20\]](#) corresponds to one LCLK DPM state as follows.

Table 70: Register Mapping for D0F0xBC_x1F2[F0:10:step20]

Register	Function	Register	Function
D0F0xBC_x1F210	State 0	D0F0xBC_x1F290	State 4
D0F0xBC_x1F230	State 1	D0F0xBC_x1F2B0	State 5
D0F0xBC_x1F250	State 2	D0F0xBC_x1F2D0	State 6
D0F0xBC_x1F270	State 3	D0F0xBC_x1F2F0	State 7

Bits	Description
31:8	Reserved.
7:0	ActivityThreshold . Read-write.

D0F0xBC_x1F300 SMU_LCLK_DPM_CNTL

Reset: xxxx_xxxxh.

Bits	Description
31:25	Reserved.
24	VoltageChgEn . Read-write. 1=Enable voltage change during LCLK DPM state transition.
23:16	LclkDpmBootState . Read-write.
15:9	Reserved.
8	LclkDpmType . Read-write.
7:1	Reserved.
0	LclkDpmEn . Read-write. 1=Enable LCLK DPM.

D0F0xBC_x1F308 SMU_LCLK_DPM_THERMAL_THROTTLING_CNTL

Reset: xxxx_xxxxh.

Bits	Description
31:25	Reserved.
24	TtHtcActive . Read-write.
23:19	Reserved.
18:16	LclkTtMode . Read-write.
15:9	Reserved.
8	TemperatureSel . Read-write.
7:1	Reserved.
0	LclkThermalThrottlingEn . Read-write.

D0F0xBC_x1F30C SMU_LCLK_DPM_THERMAL_THROTTLING_THRESHOLDS

Reset: xxxx_xxxxh.

Bits	Description
31:16	HighThreshold. Read-write. Specifies the high thermal threshold for LCLK thermal throttling.
15:0	LowThreshold. Read-write. Specifies the low thermal threshold for LCLK thermal throttling.

D0F0xBC_x1F380 FIRMWARE_FLAGS

Reset: xxxx_xxxxh.

Bits	Description						
31:24	TestCount. Read-write. Test count.						
23:1	Reserved. Read-write.						
0	InterruptsEnabled. Read-write. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Firmware has not yet enabled interrupts. BIOS/Driver cannot yet send message interrupts to SMC.</td> </tr> <tr> <td>1</td> <td>Firmware has enabled interrupts. BIOS/Driver can send message interrupts to SMC.</td> </tr> </tbody> </table>	Bits	Definition	0	Firmware has not yet enabled interrupts. BIOS/Driver cannot yet send message interrupts to SMC.	1	Firmware has enabled interrupts. BIOS/Driver can send message interrupts to SMC.
Bits	Definition						
0	Firmware has not yet enabled interrupts. BIOS/Driver cannot yet send message interrupts to SMC.						
1	Firmware has enabled interrupts. BIOS/Driver can send message interrupts to SMC.						

D0F0xBC_x1F384 FIRMWARE_VID

Reset: xxxx_xxxxh.

Bits	Description
31:8	Reserved.
7:0	FirmwareVid. Read-write. Current voltage set by firmware voltage controller.

D0F0xBC_x1F388 TEMPERATURE_READ_ADDR

Reset: xxxx_xxxxh.

Bits	Description
31:10	Reserved.
9:6	TcenId. Read-write.
5:0	CsrAddr. Read-write.

D0F0xBC_x1F39C PCIE_PG_ARGS

Read-write. Reset: xxxx_xxxxh.

Bits	Description
31:24	UpperLaneID. 0=Upper Lane ID. PHY Lanes UpperLaneID:LowerLaneID (inclusive) are affected.
23:16	LowerLaneID. 0=Lower Lane ID. PHY Lanes UpperLaneID:LowerLaneID (inclusive) are affected.
15:5	Reserved.

4	SkipCore.	<u>Definition</u>
	<u>Bits</u>	
	0	Do not skip Core Power Gating actions.
3	SkipPhy.	<u>Definition</u>
	<u>Bits</u>	
	0	Do not skip PHY Lane Power Gating actions.
2	Core.	<u>Definition</u>
	<u>Bits</u>	
	0	Do not record in Housekeeping Structure. Will NOT result in PCIe x16 Core Power Up/Down
1	Tx.	<u>Definition</u>
	<u>Bits</u>	
	0	Do not take action for PHY Tx Lanes
0	Rx. Read-write.	<u>Definition</u>
	<u>Bits</u>	
	0	Do not take action for PHY Rx Lanes

D0F0xBC_x1F3D8 LOAD_LINE_TRIM_TABLE1

Reset: xxxx_xxxxh.

Bits	Description
31:24	LoadLineTrim0. Read-write.
23:16	LoadLineTrim1. Read-write.
15:8	LoadLineTrim2. Read-write.
7:0	LoadLineTrim3. Read-write.

D0F0xBC_x1F3DC LOAD_LINE_TRIM_TABLE2

Reset: xxxx_xxxxh.

Bits	Description
31:24	LoadLineTrim4. Read-write.
23:16	LoadLineTrim5. Read-write.
15:8	LoadLineTrim6. Read-write.
7:0	LoadLineTrim7. Read-write.

D0F0xBC_x1F3F8 CSR_GNB_1

Reset: xxxx_xxxxh.

Bits	Description
31:24	SviTrimValueVddNB . Read-write. BIOS: D0F0xBC_xE0104184 [SviLoadLineTrimVddNb]. Loadline trim to use for VDDNB. Used to index into D0F0xBC_x1F3D8 and D0F0xBC_x1F3DC . See D18F5x188 [NbLoadLineTrim].
23:16	SviTrimValueVdd . Read-write. BIOS: D0F0xBC_xE0104184 [SviLoadLineTrimVdd]. Loadline trim to use for VDD. Used to index into D0F0xBC_x1F3D8 and D0F0xBC_x1F3DC . See D18F5x12C [CoreLoadLineTrim].
15:8	SviInitLoadLineVddNB . Read-write.
7:0	SviInitLoadLineVdd . Read-write.

D0F0xBC_x1F3FC CSR_GNB_2

Reset: xxxx_xxxxh.

Bits	Description
31:16	SviVidStep . Read-write.
15:0	SviVidStepBase . Read-write.

D0F0xBC_x1F400 CSR_GNB_3

Reset: xxxx_xxxxh.

Bits	Description
31:24	Reserved. Read-write.
23:16	PstateMax . Read-write. Highest enabled Pstate. This is the lowest power enabled pstate.
15:8	SviLoadLineOffsetVddNB . Read-write. Loadline offset to use for VddNB. Used to index into D0F0xBC_x1F404 . See D18F5x188 [NbOffsetTrim].
7:0	SviLoadLineOffsetVdd . Read-write. BIOS: D0F0xBC_xE0104184 [SviLoadLineOffsetVdd] Loadline offset to use for Vdd. Used to index into D0F0xBC_x1F404 . See D18F5x12C [CoreOffsetTrim].

D0F0xBC_x1F404 LOAD_LINE_OFFSET_TABLE

Reset: xxxx_xxxxh.

Bits	Description
31:24	LoadLineOffset0 . Read-write.
23:16	LoadLineOffset1 . Read-write.
15:8	LoadLineOffset2 . Read-write.
7:0	LoadLineOffset3 . Read-write.

D0F0xBC_x1F428 PM_CONFIG

Reset: xxxx_xxxxh.

Bits	Description						
31:30	Reserved. Read-write.						
29	<p>SviMode. Read-write.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>SVI1.</td> </tr> <tr> <td>1</td> <td>SVI2.</td> </tr> </tbody> </table>	Bits	Definition	0	SVI1.	1	SVI2.
Bits	Definition						
0	SVI1.						
1	SVI2.						
28	<p>BapmCoeffOverride. Read-write.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Calculate filter coefficients.</td> </tr> <tr> <td>1</td> <td>Use SW programmed filter coefficients; ignored if D18F4x15C[Boost-Lock]=1.</td> </tr> </tbody> </table>	Bits	Definition	0	Calculate filter coefficients.	1	Use SW programmed filter coefficients; ignored if D18F4x15C [Boost-Lock]=1.
Bits	Definition						
0	Calculate filter coefficients.						
1	Use SW programmed filter coefficients; ignored if D18F4x15C [Boost-Lock]=1.						
27	<p>NbPstateAllCpusIdle. Read-write.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Use low NB P-state voltage when AllCpusIdle.</td> </tr> <tr> <td>1</td> <td>Use high NB P-state voltage when AllCpusIdle.</td> </tr> </tbody> </table>	Bits	Definition	0	Use low NB P-state voltage when AllCpusIdle.	1	Use high NB P-state voltage when AllCpusIdle.
Bits	Definition						
0	Use low NB P-state voltage when AllCpusIdle.						
1	Use high NB P-state voltage when AllCpusIdle.						
26:24	PstateAllCpusIdle. Read-write. This field specifies the core P-state to use for IDD calculation when AllCpusIdle.						
23:21	Reserved.						
20	<p>IF (Revision == RL-A1) THEN</p> <p>HybridBoostEn. Read-write. This bit controls the Hybrid Boost feature. 2.5.9.1 [Hybrid Boost]. SMU initializes this bit to 1. Also see D0F0xBC_x1F8EC [HybridBoostNotSupported].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Disable Hybrid Boost.</td> </tr> <tr> <td>1</td> <td>Enable Hybrid Boost.</td> </tr> </tbody> </table> <p>ELSE Reserved. ENDIF</p>	Bits	Definition	0	Disable Hybrid Boost.	1	Enable Hybrid Boost.
Bits	Definition						
0	Disable Hybrid Boost.						
1	Enable Hybrid Boost.						
19:6	Reserved.						
5	<p>EnableNbDpm. Read-write.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Disable Dynamic NB Pstate Management. Should be cleared before sending <code>SMC_MSG_CONFIG_NBDPM</code> message to SMU.</td> </tr> <tr> <td>1</td> <td>Enable Dynamic NB Pstate Management. Should be set before sending <code>SMC_MSG_CONFIG_NBDPM</code> message to SMU.</td> </tr> </tbody> </table>	Bits	Definition	0	Disable Dynamic NB Pstate Management. Should be cleared before sending <code>SMC_MSG_CONFIG_NBDPM</code> message to SMU.	1	Enable Dynamic NB Pstate Management. Should be set before sending <code>SMC_MSG_CONFIG_NBDPM</code> message to SMU.
Bits	Definition						
0	Disable Dynamic NB Pstate Management. Should be cleared before sending <code>SMC_MSG_CONFIG_NBDPM</code> message to SMU.						
1	Enable Dynamic NB Pstate Management. Should be set before sending <code>SMC_MSG_CONFIG_NBDPM</code> message to SMU.						
4	Reserved.						
3	<p>EnableLpmx. Read-write.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Disable LPMx. Should be cleared before sending <code>SMC_MSG_CONFIG_LPMx</code> message to SMU.</td> </tr> <tr> <td>1</td> <td>Enable LPMx. Should be set before sending <code>SMC_MSG_CONFIG_LPMx</code> message to SMU.</td> </tr> </tbody> </table>	Bits	Definition	0	Disable LPMx. Should be cleared before sending <code>SMC_MSG_CONFIG_LPMx</code> message to SMU.	1	Enable LPMx. Should be set before sending <code>SMC_MSG_CONFIG_LPMx</code> message to SMU.
Bits	Definition						
0	Disable LPMx. Should be cleared before sending <code>SMC_MSG_CONFIG_LPMx</code> message to SMU.						
1	Enable LPMx. Should be set before sending <code>SMC_MSG_CONFIG_LPMx</code> message to SMU.						
2	EnableTdcLimit. Read-write.						
1	EnableBapm. Read-write.						
0	EnableVpcAccumulators. Read-write.						

D0F0xBC_x1F460 PM_INTERVAL_CNTL_0

Reset: xxxx_xxxxh.

Bits	Description
31:24	Loadline. Read-write.
23:16	VoltageCntl. Read-write.
15:8	ThermalCntl. Read-write.
7:0	LclkDpm. Read-write.

D0F0xBC_x1F468 PM_TIMER_PERIOD

Bits	Description
31:0	TimerPeriod. Read-write. Reset: X. Specifies the period at which various power management related algorithms are run. Period = TimerPeriod / REFCLK.

D0F0xBC_x1F46C PM_TIMERS_1

Bits	Description
31:24	LpmxPeriod. Read-write. Reset: 0.
23:16	BapmPeriod. Read-write. Reset: 0.
15:0	VpcPeriod. Read-write. Reset: 0.

D0F0xBC_x1F5F8 NB_PSTATE_CONFIG

Reset: xxxx_xxxxh.

Bits	Description
31:24	Reserved.
23	EnableDpmPstatePoll. Read-write
22	EnableNbPsi1. Read-write. Specifies how PSI1_L functions for VDDNB. 0=PSI1_L is deasserted. 1=PSI1_L is asserted whenever the GPU is idle.
21:18	Reserved.
17	SkipDPM0. Read-write.
16	SkipPG. Read-write.
15:8	Hysteresis. Read-write. Specifies the time the GPU must be idle before transitioning to the NB P-states indexed by Dpm0PgNbPsHi and Dpm0PgNbPsLo. Time = Hysteresis * D0F0xBC_x1F638 [NbDpmPeriod] time.
7:6	DpmXNbPsHi. Read-write. BIOS: D0F0xBC_xE010703C [NbPstateHi]. See: Dpm0PgNbPsLo.
5:4	DpmXNbPsLo. Read-write. BIOS: D0F0xBC_xE010703C [NbPstateLo]. See: Dpm0PgNbPsLo.

3:2	Dpm0PgNbPsHi. Read-write. BIOS: D0F0xBC_xE010703C[NbPstateHi] . See: Dpm0PgNbPsLo .										
1:0	Dpm0PgNbPsLo. Read-write. BIOS: D0F0xBC_xE010703C[NbPstateLo] . Indexes the NB P-state used during specific levels of GPU activity. See 2.5.4.1 [NB P-states] . <table border="1"> <thead> <tr> <th>Bits</th> <th>NB P-state Indexed</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>D18F3x160 (see D18F5x1[6C:60]).</td> </tr> <tr> <td>01b</td> <td>D18F3x164 (see D18F5x1[6C:60]).</td> </tr> <tr> <td>10b</td> <td>D18F3x168 (see D18F5x1[6C:60]).</td> </tr> <tr> <td>11b</td> <td>D18F3x16C (see D18F5x1[6C:60]).</td> </tr> </tbody> </table>	Bits	NB P-state Indexed	00b	D18F3x160 (see D18F5x1[6C:60]).	01b	D18F3x164 (see D18F5x1[6C:60]).	10b	D18F3x168 (see D18F5x1[6C:60]).	11b	D18F3x16C (see D18F5x1[6C:60]).
Bits	NB P-state Indexed										
00b	D18F3x160 (see D18F5x1[6C:60]).										
01b	D18F3x164 (see D18F5x1[6C:60]).										
10b	D18F3x168 (see D18F5x1[6C:60]).										
11b	D18F3x16C (see D18F5x1[6C:60]).										

D0F0xBC_x1F5FC NB_PSTATE_STATUS

Reset: xxxx_xxxxh.

Bits	Description
31:9	Reserved.
8	PSI1Sts. Read-write. 1=PSI1_L is enabled for VDDNB. 0=PSI1_L is disabled for VDDNB.
7:2	Reserved.
1	CurrentPstatePair. Read-write, updated-by-hardware. Specifies the NB P-state pair in use. 1= D0F0xBC_x1F5F8[Dpm0PgNbPsHi, Dpm0PgNbPsLo] . 0= D0F0xBC_x1F5F8[DpmXNbPsHi, DpmXNbPsLo] .
0	ChangeInProgress. Read-write, updated-by-hardware. 1=Hardware is updating D18F5x170[NbPstateHi, NbPstateLo] .

D0F0xBC_x1F628 LHTC CONFIG 1

See [2.10.4.2 \[Local Hardware Thermal Control \(LHTC\)\]](#).

Bits	Description
31:24	Reserved.
23:16	LhtcActivePstateLimit. Read-write. Reset: 0. BIOS: D0F0xBC_xE0104188[LhtcPstateLimit] . Specifies the P-state limit of all cores when in the LHTC-active state. LhtcPstateLimit must be greater than or equal to its reset value (same or lower performing P-state) but not exceed MSRC001_0061[PstateMaxVal] + D18F4x15C[NumBoostStates] . This field uses hardware P-state numbering, see 2.5.3.1.2.2 [Hardware P-state Numbering] .
15:0	Reserved.

D0F0xBC_x1F62C TDC_VRM_LIMIT

Bits	Description
31:16	Iddnb. Read-write. Reset: 0.
15:0	Idd. Read-write. Reset: 0.

D0F0xBC_x1F638 PM_TIMERS_2

Reset: xxxx_xxxxh.

Bits	Description
31:24	PgInterlockPeriod. Read-write.
23:16	NbdpmPeriod. Read-write. Specifies the rate at which hardware determines which two NB P-states should be in use. Time = NbdpmPeriod * D0F0xBC_x1F468[TimerPeriod] time.
15:8	HtcPeriod. Read-write.
7:0	TdcPeriod. Read-write.

D0F0xBC_x1F86C LHTC CONFIG 2

Reset: xxxx_xxxxh.

Bits	Description
31:24	Reserved.
23	LhtcCap. IF D18F4x15C[BoostLock] THEN Read-only, updated-by-SMU. ELSE Read-write, updated-by-SMU. ENDIF. 1=D0F0xBC_x1F8555[LhtcPstateLimit] is valid. 0=D0F0xBC_x1F8555[LhtcPstateLimit] is not valid. See 2.10.4.2 [Local Hardware Thermal Control (LHTC)].
22:0	Reserved.

D0F0xBC_x1F89C LHTC CONFIG 3

Reset: xxxx_xxxxh.

Bits	Description								
31:24	LhtcGtempLimitLo. IF D18F4x15C[BoostLock] THEN Read-only, updated-by-SMU. ELSE Read-write, updated-by-SMU. ENDIF. See 2.10.4.2 [Local Hardware Thermal Control (LHTC)]. See HtcGtempLimitHi for encoding.								
23:16	LhtcGtempLimitHi. IF D18F4x15C[BoostLock] THEN Read-only, updated-by-SMU. ELSE Read-write, updated-by-SMU. ENDIF. See 2.10.4.2 [Local Hardware Thermal Control (LHTC)]. Range is -49°C to 206°C and is encoded as <Limit-49>°C. <table border="1"> <thead> <tr> <th>Bits</th> <th>Temperature Limit</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>-49°C.</td> </tr> <tr> <td>FEh-01h</td> <td><Limit - 49> °C.</td> </tr> <tr> <td>FFh</td> <td>206°C.</td> </tr> </tbody> </table>	Bits	Temperature Limit	00h	-49°C.	FEh-01h	<Limit - 49> °C.	FFh	206°C.
Bits	Temperature Limit								
00h	-49°C.								
FEh-01h	<Limit - 49> °C.								
FFh	206°C.								
15:0	Reserved.								

D0F0xBC_x1F8EC HYBRID_BOOST_CONFIG

Reset: xxxx_xxxxh.

IF (Revision==RL-A1) THEN

Bits	Description
31:25	Reserved.

24	HybridBoostNotSupported. IF D18F4x15C [BoostLock] THEN Read-only, updated-by-SMU. ELSE Read-write, updated-by-SMU. ENDIF. Hybrid Boost cannot be enabled if this bit is set.
23:0	Reserved.

ELSE

Bits	Description
31:0	Reserved.

ENDIF.

D0F0xBC_xE000120 Activity Monitor Control

Bits	Description												
31:11	Reserved.												
10	EnOrbDsCnt. Read-write. Reset: 0. 1=Enable downstream counter.												
9	EnOrbUsCnt. Read-write. Reset: 0. 1=Enable upstream counter.												
8	EnBifCnt. Read-write. Reset: 0. 1=Enable BIF counter.												
7:5	Reserved.												
4:3	BusyCntSel. Read-write. Reset: 0. Specifies subcomponents or activity monitored by the LCLK activity monitor. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>GFX DMA (BIF)</td> <td>10b</td> <td>Downstream activity</td> </tr> <tr> <td>01b</td> <td>Upstream activity</td> <td>11b</td> <td>Up/downstream activity max</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	00b	GFX DMA (BIF)	10b	Downstream activity	01b	Upstream activity	11b	Up/downstream activity max
Bits	Definition	Bits	Definition										
00b	GFX DMA (BIF)	10b	Downstream activity										
01b	Upstream activity	11b	Up/downstream activity max										
2	Reserved.												
1	PeriodCntRst. Read-write. Reset: 1.												
0	ActivityCntRst. Read-write. Reset: 1.												

D0F0xBC_xE0003000 CPU Interrupt Request

See 2.13.1 [Software Interrupts].

Bits	Description
31:17	Reserved.
16:1	ServiceIndex. Read-write. Reset: 0.
0	IntToggle. Read-write. Reset: 0.

D0F0xBC_xE0003004 CPU Interrupt Status

See 2.13.1 [Software Interrupts].

Bits	Description
31:2	Reserved.
1	IntDone. Read-only; updated-by-hardware. Reset: 0.
0	IntAck. Read-only; updated-by-hardware. Reset: 0.

D0F0xBC_xE0003048 SCLK_MIN_DIV

Bits	Description
31:19	Reserved.
18:12	Intv. Read-write. Reset: 02h. Divider integer value
11:0	Fracv. Read-write. Reset: 0. Divider fraction value

D0F0xBC_xE0003088 SMU_AUTH_STATUS

Bits	Description
31:2	Reserved.
1	SmuAuthPass. Read-write. Reset: 0.
0	SmuAuthDone. Read-write. Reset: 0.

D0F0xBC_xE00030A4 SMU_FIRMWARE_AUTH

Bits	Description
31:17	Reserved.
16	SmuProtectedMode. Read-only. Reset: 0.
15:0	Reserved.

D0F0xBC_xE0104168 MEMCLK VID Configuration

Bits	Description
31:30	MemClkVid3[1:0]. Read-only. Reset: value varies by product. See MemClkVid0
29:22	MemClkVid2. Read-only. Reset: value varies by product. See MemClkVid0.
21:14	MemClkVid1. Read-only. Reset: value varies by product. See MemClkVid0.
13:6	<p>MemClkVid0. Read-only. Reset: value varies by product. Specifies the VDDNB voltage required for a specific MEMCLK. See the <i>AMD Serial VID Interface 2.0 (SVI2) Specification</i>.</p> <p>Each MemClkVid field corresponds to a MEMCLK frequency as follows:</p> <ul style="list-style-type: none"> • MemClkVid8: 1200MHz. • MemClkVid7: 1066MHz. • MemClkVid6: 1050MHz. • MemClkVid5: 933MHz. • MemClkVid4: 800MHz. • MemClkVid3: 667MHz. • MemClkVid2: 533MHz. • MemClkVid1: 400MHz. • MemClkVid0: 333 MHz.
5:0	Reserved.

D0F0xBC_xE010416C MEMCLK VID Configuration

Bits	Description
31:30	MemClkVid7[1:0] . Read-only. Reset: value varies by product. See D0F0xBC_xE0104168[MemClkVid0] .
29:22	MemClkVid6 . Read-only. Reset: value varies by product. See D0F0xBC_xE0104168[MemClkVid0] .
21:14	MemClkVid5 . Read-only. Reset: value varies by product. See D0F0xBC_xE0104168[MemClkVid0] .
13:6	MemClkVid4 . Read-only. Reset: value varies by product. See D0F0xBC_xE0104168[MemClkVid0] .
5:0	MemClkVid3[7:2] . Read-only. Reset: value varies by product. See D0F0xBC_xE0104168[MemClkVid0] .

D0F0xBC_xE0104170 MEMCLK VID Configuration

Bits	Description
31:14	Reserved.
13:6	MemClkVid8 . Read-only. Reset: value varies by product. See D0F0xBC_xE0104168[MemClkVid0] .
5:0	MemClkVid7[7:2] . Read-only. Reset: value varies by product. See D0F0xBC_xE0104168[MemClkVid0] .

D0F0xBC_xE0104184 SVI Loadline Configuration

Bits	Description
31:10	Reserved.
9:8	SviLoadLineOffsetVddNb . Read-only. Reset: value varies by product. See D18F5x188[NbOffsetTrim] .
7:6	SviLoadLineOffsetVdd . Read-only. Reset: value varies by product. See D18F5x12C[CoreOffsetTrim] .
5:3	SviLoadLineTrimVddNb . Read-only. Reset: value varies by product. See D18F5x188[NbLoadLineTrim] .
2:0	SviLoadLineTrimVdd . Read-only. Reset: value varies by product. See D18F5x12C[CoreLoadLineTrim] .

D0F0xBC_xE0104188 SMU Power Management Config 1

Bits	Description
31	Reserved.
30	BapmDisable . Read-only.
29	BapmMeasuredTemp . Read-only.

28:26	LhtcPstateLimit. Read-only. Reset: value varies by product. See D0F0xBC_x1F628 [LhtcPstateLimit] and 2.10.4.2 [Local Hardware Thermal Control (LHTC)].
25:0	Reserved.

D0F0xBC_xE010418C SMU Power Management Config 2

Bits	Description
31:0	Reserved.

D0F0xBC_xE010703C SMU Power Management Config 3

Bits	Description
31:7	Reserved.
6:5	NbPstateLo. Read-only. Reset: value varies by product. See D0F0xBC_x1F5F8 [Dpm0PgNbPsLo].
4:3	NbPstateHi. Read-only. Reset: value varies by product. See D0F0xBC_x1F5F8 [Dpm0PgNbPsHi].
2:0	Reserved.

D0F0xBC_xE0300000 SBLK_MC_PGFSM_CONFIG

Read-write. Reset: 0.

Bits	Description
31:28	RegAddr.
27:14	Reserved.
13	ReadOp.
12	WriteOp.
11	P2Select.
10	P1Select.
9	PowerUp.
8	PowerDown.
7:0	FsmAddr.

D0F0xBC_xE030000C SBLK_PGD_PGFSM_CONFIG

Bits	Description
31:28	RegAddr. Read-write. Reset: 0.
27:14	Reserved.
13	ReadOp. Read-write. Reset: 0.
12	WriteOp. Read-write. Reset: 0.

11	P2Select. Read-write. Reset: 0.
10	P1Select. Read-write. Reset: 0.
9	PowerUp. Read-write. Reset: 0.
8	PowerDown. Read-write. Reset: 0.
7:0	FsmAddr. Read-write. Reset: 0.

D0F0xBC_xE0300018 SBLK_IOMMU_PGFSM_CONFIG

Bits	Description
31:28	RegAddr. Read-write. Reset: 0.
27:14	Reserved.
13	ReadOp. Read-write. Reset: 0.
12	WriteOp. Read-write. Reset: 0.
11	P2Select. Read-write. Reset: 0.
10	P1Select. Read-write. Reset: 0.
9	PowerUp. Read-write. Reset: 0.
8	PowerDown. Read-write. Reset: 0.
7:0	FsmAddr. Read-write. Reset: 0.

D0F0xBC_xE0300024 SBLK_VCE_PGFSM_CONFIG

Bits	Description
31:28	RegAddr. Read-write. Reset: 0.
27:14	Reserved.
13	ReadOp. Read-write. Reset: 0.
12	WriteOp. Read-write. Reset: 0.
11	P2Select. Read-write. Reset: 0.
10	P1Select. Read-write. Reset: 0.
9	PowerUp. Read-write. Reset: 0.
8	PowerDown. Read-write. Reset: 0.
7:0	FsmAddr. Read-write. Reset: 0.

D0F0xBC_xE0300030 SBLK_DC2_PGFSM_CONFIG

Bits	Description
31:28	RegAddr. Read-write. Reset: 0.
27:14	Reserved.
13	ReadOp. Read-write. Reset: 0.
12	WriteOp. Read-write. Reset: 0.

11	P2Select. Read-write. Reset: 0.
10	P1Select. Read-write. Reset: 0.
9	PowerUp. Read-write. Reset: 0.
8	PowerDown. Read-write. Reset: 0.
7:0	FsmAddr. Read-write. Reset: 0.

D0F0xBC_xE030003C UVD_CHAIN_PGFSM_CONFIG

Bits	Description
31:28	RegAddr. Read-write. Reset: 0.
27:14	Reserved.
13	ReadOp. Read-write. Reset: 0.
12	WriteOp. Read-write. Reset: 0.
11	P2Select. Read-write. Reset: 0.
10	P1Select. Read-write. Reset: 0.
9	PowerUp. Read-write. Reset: 0.
8	PowerDown. Read-write. Reset: 0.
7:0	FsmAddr. Read-write. Reset: 0.

D0F0xBC_xE0300200 SBLK_MC_PGFSM_DEBUG

Bits	Description
31:11	Reserved.
10	P1IsoN. Read-only. Reset: 1. P1 isolation
9:0	Reserved.

D0F0xBC_xE0300208 SBLK_IOMMU_PGFSM_DEBUG

Bits	Description
31:11	Reserved.
10	P1IsoN. Read-only. Reset: 1. P1 isolation
9:0	Reserved.

D0F0xBC_xE030020C SBLK_VCE_PGFSM_DEBUG

Bits	Description
31:11	Reserved.
10	P1IsoN. Read-only. Reset: 1. P1 isolation
9:0	Reserved.

D0F0xBC_xE0300210 SBLK_DC2_PGFSM_DEBUG

Bits	Description
31:11	Reserved.
10	P1IsoN . Read-only. Reset: 1. P1 isolation
9:0	Reserved.

D0F0xBC_xE0300218 SBLK_UVDU_PGFSM_DEBUG

Bits	Description
31:11	Reserved.
10	P1IsoN . Read-only. Reset: 1. P1 isolation
9:0	Reserved.

D0F0xBC_xE0300320 PGFSM_CLK_CNTL_0

Bits	Description
31:2	Reserved.
1	IommuPgfsmClockEn . Read-write. Reset: 1.
0	PgdPgfsmClockEn . Read-write. Reset: 1.

D0F0xBC_xE0300324 PGFSM_CLK_CNTL_1

Bits	Description
31:3	Reserved.
2	Dc2PgfsmClockEn . Read-write. Reset: 1.
1	UvdPgfsmClockEn . Read-write. Reset: 1.
0	VcePgfsmClockEn . Read-write. Reset: 1.

D0F0xE0 Link Index Address

Reset: 0130_8001h.

[D0F0xE0](#) and [D0F0xE4](#) are used to access [D0F0xE4_x\[FFFF_FFFF:0000_0000\]](#). To read or write to one of these register, the address is written first into the address register [D0F0xE0](#) and then the data is read from or written to the data register [D0F0xE4](#).

The phy index registers ([D0F0xE4_x\[2xxx_xxxx\]](#)) mapping to a specific phy, pin or pin group is shown in a table in the register definition. For example, to perform a read or write operation to configure Gfx phy 0 (P_GFX_[T,R]X[P,N][7:0] pin group) compensation, software should program [D0F0xE0\[31:0\]=0121_0000h](#). Accessing any register number that is not listed in the mapping table may result in undefined behavior.

Some phy registers support broadcast write operations to groups of 4 or 8 lanes. For example, to perform

broadcast write operation to configure Gfx Link[3:0] (P_GFX_RX[P,N][3:0] lanes) receiver phase loop filter, software should program D0F0xE0[31:0]=0221_5602h.

Bits	Description																				
31:24	BlockSelect: block select. Read-write. This field is used to select the specific register block to access. The encodings supported depends on the FrameType selected. <table border="1"> <thead> <tr> <th>FrameType</th> <th>Encoding</th> </tr> </thead> <tbody> <tr> <td>01h</td> <td>1=GPP link core, 2=Gfx link core</td> </tr> <tr> <td>1xh</td> <td>1=Phy interface 0, 2=Phy interface 1 (FrameType 11h only)</td> </tr> <tr> <td>2xh</td> <td>1=Phy 0, 2=Phy 1 (FrameType 21h only)</td> </tr> <tr> <td>3xh</td> <td>1=Wrapper</td> </tr> </tbody> </table>	FrameType	Encoding	01h	1=GPP link core, 2=Gfx link core	1xh	1=Phy interface 0, 2=Phy interface 1 (FrameType 11h only)	2xh	1=Phy 0, 2=Phy 1 (FrameType 21h only)	3xh	1=Wrapper										
FrameType	Encoding																				
01h	1=GPP link core, 2=Gfx link core																				
1xh	1=Phy interface 0, 2=Phy interface 1 (FrameType 11h only)																				
2xh	1=Phy 0, 2=Phy 1 (FrameType 21h only)																				
3xh	1=Wrapper																				
23:16	FrameType: frame type. Read-write. This field is used to select the type of register block to access. <table border="1"> <thead> <tr> <th>Bits</th> <th>Destination</th> </tr> </thead> <tbody> <tr> <td>01h</td> <td>Link core registers</td> </tr> <tr> <td>1Nh</td> <td>Phy interface block registers.</td> </tr> <tr> <td>2Nh</td> <td>Phy registers.</td> </tr> <tr> <td>3Nh</td> <td>Wrapper registers.</td> </tr> </tbody> </table> <table border="1"> <thead> <tr> <th>N</th> <th>Register Block</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>GPP PCIe Links</td> </tr> <tr> <td>1h</td> <td>Gfx PCIe links</td> </tr> <tr> <td>2h</td> <td>DDI</td> </tr> <tr> <td>3h</td> <td>DDI2</td> </tr> </tbody> </table>	Bits	Destination	01h	Link core registers	1Nh	Phy interface block registers.	2Nh	Phy registers.	3Nh	Wrapper registers.	N	Register Block	0h	GPP PCIe Links	1h	Gfx PCIe links	2h	DDI	3h	DDI2
Bits	Destination																				
01h	Link core registers																				
1Nh	Phy interface block registers.																				
2Nh	Phy registers.																				
3Nh	Wrapper registers.																				
N	Register Block																				
0h	GPP PCIe Links																				
1h	Gfx PCIe links																				
2h	DDI																				
3h	DDI2																				
15:0	PcieIndxAddr: index address. Read-write.																				

D0F0xE4 Link Index Data

See D0F0xE0.

Bits	Description
31:0	PcieIndxData: index data. Read-write.

3.3.1 IO Link Registers

D0F0xE4_x0[2:1]01_0002 IO Link Hardware Debug

Reset: 0000_0000h.

Bits	Description
31:1	Reserved.
0	HwDebug[0]: ignore DLLPs in L1. Read-write. BIOS: 1. 1=DLLPs are ignored in L1 so the TXCLK can be turned off.

D0F0xE4_x0[2:1]01_0010 IO Link Control 1

Reset: 80E3_0800h.

Bits	Description
31:13	Reserved.

12:10	RxSbAdjPayloadSize. Read-write. BIOS: 100b. Payload size for DMA requests from the PCIe controller buffer.																
	<table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>00xb</td> <td>Reserved.</td> <td>100b</td> <td>64 bytes.</td> </tr> <tr> <td>010b</td> <td>16 bytes.</td> <td>101b</td> <td>Reserved.</td> </tr> <tr> <td>011b</td> <td>32 bytes.</td> <td>11xb</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	00xb	Reserved.	100b	64 bytes.	010b	16 bytes.	101b	Reserved.	011b	32 bytes.	11xb	Reserved.
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>														
00xb	Reserved.	100b	64 bytes.														
010b	16 bytes.	101b	Reserved.														
011b	32 bytes.	11xb	Reserved.														
9	UmiNpMemWrite: memory write mapping enable. Read-write. 1=Internal non-posted memory writes are transferred to UMI.																
8:1	Reserved.																
0	HwInitWrLock: hardware init write lock. Read-write. 1=Lock HWInit registers. 0=Unlock HWInit registers.																

D0F0xE4_x0[2:1]01_0011 IO Link Config Control

Reset: 0000_0007h.

Bits	Description																				
31:30	Reserved.																				
29:27	CiPrivMaxCplPayloadSize. Read-write. Private value for MAX_PAYLOAD_SIZE, used when CiMaxCplPayloadSizeMode=1.																				
26	CiMaxCplPayloadSizeMode. Read-write. Used for Slave completions. Private override for cfg register DEVICE_CNTL:MAX_PAYLOAD_SIZE. 1=Use CiPrivMaxCplPayloadSize.																				
25	CiExtendedTagEnOverride. Read-write. 1=Private override for DEVICE_CNTL:EXTENDED_TAG_EN.																				
24	CiMaxReadSafeMode. Read-write. 1=Only 32 or 64 byte address aligned reads are generated.																				
23:21	CiPrivMaxReadRequestSize. Read-write. Specifies the maximum read request size when CiMaxReadRequestSizeMode=1. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>128 bytes.</td> <td>100b</td> <td>2 Kbytes.</td> </tr> <tr> <td>001b</td> <td>256 bytes.</td> <td>101b</td> <td>4 Kbytes.</td> </tr> <tr> <td>010b</td> <td>512 bytes.</td> <td>110b</td> <td>Reserved.</td> </tr> <tr> <td>011b</td> <td>1 Kbyte.</td> <td>111b</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	000b	128 bytes.	100b	2 Kbytes.	001b	256 bytes.	101b	4 Kbytes.	010b	512 bytes.	110b	Reserved.	011b	1 Kbyte.	111b	Reserved.
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>																		
000b	128 bytes.	100b	2 Kbytes.																		
001b	256 bytes.	101b	4 Kbytes.																		
010b	512 bytes.	110b	Reserved.																		
011b	1 Kbyte.	111b	Reserved.																		
20	CiMaxReadRequestSizeMode. Read-write. 1=CiPrivMaxReadRequestSize defines the maximum read request size. 0=D[8:2]F0x60[MaxRequestSize] defines the maximum read request size.																				
19:17	CiPrivMaxPayloadSize. Read-write. Specifies the maximum payload size when CiMaxPayloadSizeMode=1. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>128 bytes.</td> <td>100b</td> <td>2 Kbytes.</td> </tr> <tr> <td>001b</td> <td>256 bytes.</td> <td>101b</td> <td>4 Kbytes.</td> </tr> <tr> <td>010b</td> <td>512 bytes.</td> <td>110b</td> <td>Reserved.</td> </tr> <tr> <td>011b</td> <td>1 Kbyte.</td> <td>111b</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	000b	128 bytes.	100b	2 Kbytes.	001b	256 bytes.	101b	4 Kbytes.	010b	512 bytes.	110b	Reserved.	011b	1 Kbyte.	111b	Reserved.
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>																		
000b	128 bytes.	100b	2 Kbytes.																		
001b	256 bytes.	101b	4 Kbytes.																		
010b	512 bytes.	110b	Reserved.																		
011b	1 Kbyte.	111b	Reserved.																		
16	CiMaxPayloadSizeMode. Read-write. 1=CiPrivMaxPayloadSize defines the maximum payload size. 0=D[8:2]F0x60[MaxPayloadSize] defines the maximum payload size.																				
15:4	Reserved.																				
3:0	DynClkLatency: dynamic clock latency. Read-write. BIOS: See 2.11.4.2.2. Specifies the number of clock cycles after logic goes idle before clocks are gated off.																				

D0F0xE4_x0[2:1]01_001C IO Link Control 2

Reset: 0E00_0000h.

Bits	Description
31:11	Reserved.
10:6	TxArbMstLimit: transmitter arbitration master limit. Read-write. BIOS: 4h. Defines together with TxArbSlvLimit a round robin arbitration pattern for downstream accesses. TxArbMstLimit defines the weight for downstream CPU requests and TxArbSlvLimit for the downstream read responses.
5:1	TxArbSlvLimit: transmitter arbitration slave limit. Read-write. BIOS: 4h. See TxArbMstLimit for details
0	TxArbRoundRobinEn: transmitter round robin arbitration enabled. Read-write. BIOS: 1. 1=Enable transmitter round robin arbitration. 0=Disable transmitter round robin arbitration.

D0F0xE4_x0[2:1]01_0020 IO Link Chip Interface Control

Reset: 0000_0050h.

Bits	Description
31:10	Reserved.
9	CiRcOrderingDis: chip interface RC ordering disable. Read-write. 0=RC ordering logic is enabled. 1=RC ordering logic is disabled.
8	CiSlvOrderingDis: slave interface ordering disable. Read-write. 0=Slave ordering logic is enabled. 1=Slave ordering logic is disabled.
7:0	Reserved.

D0F0xE4_x0[2:1]01_0040 IO Link Phy Control

Reset: 0000_0000h.

Bits	Description										
31:16	Reserved.										
15:14	PElecIdleMode: electrical idle mode for physical layer. Read-write. BIOS: 10b. Defines which electrical idle signal is used, either inferred by link controller or from phy. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Gen1 - entry:phy, exit:phy; Gen2 - entry:infer, exit:phy</td> </tr> <tr> <td>01b</td> <td>Gen1 - entry:infer, exit:phy; Gen2 - entry:infer, exit:phy</td> </tr> <tr> <td>10b</td> <td>Gen1 - entry:phy, exit:phy; Gen2 - entry:PHY, exit:phy</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Definition	00b	Gen1 - entry:phy, exit:phy; Gen2 - entry:infer, exit:phy	01b	Gen1 - entry:infer, exit:phy; Gen2 - entry:infer, exit:phy	10b	Gen1 - entry:phy, exit:phy; Gen2 - entry:PHY, exit:phy	11b	Reserved
Bits	Definition										
00b	Gen1 - entry:phy, exit:phy; Gen2 - entry:infer, exit:phy										
01b	Gen1 - entry:infer, exit:phy; Gen2 - entry:infer, exit:phy										
10b	Gen1 - entry:phy, exit:phy; Gen2 - entry:PHY, exit:phy										
11b	Reserved										
13:0	Reserved.										

D0F0xE4_x0[2:1]01_00B0 IO Link Strap Link Strap Control

Reset: 0000_0001h.

Bits	Description
31:6	Reserved.
5	StrapF0AerEn. Read-write. BIOS: 0. 1=AER support enabled. 0=AER support disabled.
4:3	Reserved.
2	StrapF0MsiEn. Read-write. BIOS: 1. 1=MSI enabled. 0=MSI disabled.
1:0	Reserved.

D0F0xE4_x0[2:1]01_00C0 IO Link Strap Miscellaneous

Bits	Description
31	Reserved.
30	StrapFlrEn. Read-write. Reset:0.
29	StrapMstAdr64En. Read-write. Reset: 0.
28	StrapReverseAll. Read-write. Reset: 0. Reverse All strap override.
27:0	Reserved.

D0F0xE4_x0[2:1]01_00C1 IO Link Strap Miscellaneous

Bits	Description
31:2	Reserved.
1	StrapGen2Compliance. Read-write. Reset: 1.
0	StrapLinkBwNotificationCapEn. Read-write. Reset: 0.

3.3.2 PIF Registers**D0F0xE4_x0[2:1]1[3:0]_0010 PIF Control**

Reset: 018A_0059h.

Table 71: Index addresses for D0F0xE4_x0[2:1]1[3:0]_0010

D0F0xE0[31:16]	D0F0xE0[15:0]
	0010h
0110h	GPPFCH PIF
0111h	Gfx PIF 0
0211h	Gfx PIF 1
0112h	DDI PIF
0113h	DDI2 PIF

Bits	Description																				
31:20	Reserved.																				
19:17	LS2ExitTime: LS2 exit time. Read-write. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>10us</td> <td>100b</td> <td>625ns</td> </tr> <tr> <td>001b</td> <td>5us</td> <td>101b</td> <td>0s</td> </tr> <tr> <td>010b</td> <td>2.5us</td> <td>110b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>1.25us</td> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	000b	10us	100b	625ns	001b	5us	101b	0s	010b	2.5us	110b	Reserved	011b	1.25us	111b	Reserved
Bits	Definition	Bits	Definition																		
000b	10us	100b	625ns																		
001b	5us	101b	0s																		
010b	2.5us	110b	Reserved																		
011b	1.25us	111b	Reserved																		
16:8	Reserved.																				
7	RxDetectTxPwrMode: receiver detection transmitter power mode. Read-write. 1=Transmitter is powered on.																				
6	RxDetectFifoResetMode: receiver detect FIFO reset mode. Read-write. BIOS: 1. 1=The transmit FIFO is reset after receiver detection. 0=The transmit FIFO is not reset after receiver detection.																				
5	Reserved.																				
4	EiDetCycleMode: electrical idle detect mode. Read-write. 1=Electrical idle cycle detection mode is enabled in L1. 0=Electrical idle detection is always enabled in L1.																				
3:0	Reserved.																				

D0F0xE4_x0[2:1]1[3:0]_0011 PIF Pairing

Reset: 0200_0000h.

Table 72: Index addresses for D0F0xE4_x0[2:1]1[3:0]_0011

D0F0xE0[31:16]	D0F0xE0[15:0]
	0011h
0110h	GPPFCH PIF
0111h	Gfx PIF 0
0211h	Gfx PIF 1
0112h	DDI PIF
0113h	DDI2 PIF

Bits	Description
31:26	Reserved.
25	MultiPif: x16 link. Read-write. 1=Lanes 7:0 are paired with a second PIF to create a x16 link.
24:17	Reserved.
16	X8Lane70: x8 link lanes 7:0. Read-write. 1=Lanes 7:0 are paired to create a x8 link.
15:13	Reserved.
12	X4Lane52: x4 link lanes 5:2. Read-write. 1=Lanes 5:2 are paired to create a x4 link.
11:10	Reserved.
9	X4Lane74: x4 link lanes 7:4. Read-write. 1=Lanes 7:4 are paired to create a x4 link.
8	X4Lane30: x4 link lanes 3:0. Read-write. 1=Lanes 3:0 are paired to create a x4 link.
7:4	Reserved.

3	X2Lane76: x2 link lanes 7:6. Read-write. 1=Lanes 7:6 are paired to create a x2 link.
2	X2Lane54: x2 link lanes 5:4. Read-write. 1=Lanes 5:4 are paired to create a x2 link.
1	X2Lane32: x2 link lanes 3:2. Read-write. 1=Lanes 3:2 are paired to create a x2 link.
0	X2Lane10: x2 link lanes 1:0. Read-write. 1=Lanes 1:0 are paired to create a x2 link.

D0F0xE4_x0[2:1]1[3:0]_001[3:2] PIF Power Down Control [1:0]

Reset: 0001_0022h.

Table 73: Index addresses for D0F0xE4_x0[2:1]1[3:0]_001[3:2]

D0F0xE0[31:16]	D0F0xE0[15:0]	
	0013h	0012h
0110h	GPPFCH PIF Lanes 7-4	GPPFCH PIF Lanes 3-0
0111h	Gfx PIF 0 Lanes 7-4	Gfx PIF 0 Lanes 3-0
0211h	Gfx PIF 1 Lanes 7-4	Gfx PIF 1 Lanes 3-0
0112h	DDI PIF Lanes 7-4	DDI PIF Lanes 3-0
0113h	DDI2 PIF Lanes 6-4	DDI2 PIF Lanes 3-0

Bits	Description																				
31:29	PIIPwrOverrideVal: PLL power state override value. Read-write. See TxPowerStateInTxs2.																				
28	PIIPwrOverrideEn: PLL power state override enable. Read-write. 1=PLL forced to the power state specified by PIIPwrOverrideVal.																				
27	Reserved.																				
26:24	PIIRampUpTime: PLL ramp time. Read-write. <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>5us</td> <td>100b</td> <td>0s</td> </tr> <tr> <td>001b</td> <td>10us</td> <td>101b</td> <td>0s</td> </tr> <tr> <td>010b</td> <td>300us</td> <td>110b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>500us</td> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	000b	5us	100b	0s	001b	10us	101b	0s	010b	300us	110b	Reserved	011b	500us	111b	Reserved
Bits	Definition	Bits	Definition																		
000b	5us	100b	0s																		
001b	10us	101b	0s																		
010b	300us	110b	Reserved																		
011b	500us	111b	Reserved																		
23:17	Reserved.																				
16	Tx2p5clkClockGatingEn. Read-write. 1=The 2.5x TxClk is gated if the lane is idle 0=The 2.5x TxClk is never gated.																				
15:13	Reserved.																				
12:10	PIIPowerStateInOff: PLL off power state. Read-write. See: TxPowerStateInTxs2. All links associated with the PLL must be in the off state to transition the PLL to this state.																				
9:7	PIIPowerStateInTxs2: PLL L1 power state. Read-write. See: TxPowerStateInTxs2. All links associated with the PLL must be in L1 to transition the PLL to this state.																				
6:4	RxPowerStateInRxs2: receiver L1 power state. Read-write. See: TxPowerStateInTxs2.																				

3	ForceRxEnInL0s: force receiver enable in L0s. Read-write. 1=The phy CDR is always enabled in L0s.			
2:0	TxPowerStateInTxs2: transmitter L1 power state. Read-write.			
	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>
	000b	L0	100b	Reserved
	001b	LS1	101b	Reserved
	010b	LS2	110b	Reserved
	011b	Reserved	111b	Off

D0F0xE4_x0[2:1]1[3:0]_0015 PIF Transmitter Status

Reset: 0000_0000h.

Table 74: Index addresses for D0F0xE4_x0[2:1]1[3:0]_0015

D0F0xE0[31:16]	D0F0xE0[15:0]
	0015h
0110h	GPPFCH PIF
0111h	Gfx PIF 0
0211h	Gfx PIF 1
0112h	DDI PIF
0113h	DDI2 PIF

Bits	Description
31:8	Reserved.
7	TxPhyStatus07: lane 7 TxPhyStatus. Read-only. See: TxPhyStatus00.
6	TxPhyStatus06: lane 6 TxPhyStatus. Read-only. See: TxPhyStatus00.
5	TxPhyStatus05: lane 5 TxPhyStatus. Read-only. See: TxPhyStatus00.
4	TxPhyStatus04: lane 4 TxPhyStatus. Read-only. See: TxPhyStatus00.
3	TxPhyStatus03: lane 3 TxPhyStatus. Read-only. See: TxPhyStatus00.
2	TxPhyStatus02: lane 2 TxPhyStatus. Read-only. See: TxPhyStatus00.
1	TxPhyStatus01: lane 1 TxPhyStatus. Read-only. See: TxPhyStatus00.
0	TxPhyStatus00: lane 0 TxPhyStatus. Read-only. Returns the state of the TxPhyStatus signal From the PIF to the BIF.

3.3.3 Phy Registers

There are three categories of phy registers: [3.3.3.1 \[Global Phy Control Registers\]](#), receiver lane control registers and transmitter lane control registers.

3.3.3.1 Global Phy Control Registers

Each global phy control register may have one instance per phy or two instances per phy (one per nibble). When a global register is implemented per phy the mapping to signal pins is shown in [Table 75](#). When a global register is implemented per nibble the mapping to pins is shown in [Table 75](#).

Table 75: Per phy register addresses to pin mappings

D0F0xE0[31:0]	Pin Names
0120h_[2:0]xxxh	GPP Links: P_GPP_[T,R]X[P,N][3:0] & FCH ports: P_UMI_[T,R]X[P,N][3:0]
0121h_[2:0]xxxh	Gfx Links[7:0]: P_GFX_[T,R]X[P,N][7:0]
0221h_[2:0]xxxh	Gfx Links[15:8]: P_GFX_[T,R]X[P,N][15:8]
0122h_[2:0]xxxh	DDI 1: DP1_TX[P,N][3:0] & DDI 0: DP0_TX[P,N][3:0]
0123h_[2:0]xxxh	DDI 2: DP2_TX[P,N][6:0]

Table 76: Per nibble register addresses to pin mappings

D0F0xE0[31:12]	Pin Names	
	Address N+1	Address N
0120h_[2:0]xxxh	GPP ports: P_GPP_[T,R]X[P,N][3:0]	FCH ports: P_UMI_[T,R]X[P,N][3:0]
0121h_[2:0]xxxh	Graphics port lower: P_GFX_[T,R]X[P,N][7:4]	Graphics port lower: P_GFX_[T,R]X[P,N][3:0]
0221h_[2:0]xxxh	Graphics port upper: P_GFX_[T,R]X[P,N][15:12]	Graphics port upper: P_GFX_[T,R]X[P,N][11:8]
0122h_[2:0]xxxh	DDI 0: DP0_TX[P,N][3:0]	DDI 1: DP1_TX[P,N][3:0]
0123h_[2:0]xxxh	DDI 2:DP2_TX[P,N][6:4]	DDI 2:DP2_TX[P,N][3:0]

D0F0xE4_x0[2:1]2[3:0]_0000 Phy Compensation Control and Calibration Control I

This register provides general control of various circuits that perform auto-calibration.

Table 77: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_0000

D0F0xE0[31:16]	D0F0xE0[15:0]
	0000h
0120h	GPP Links: P_GPP_[T,R]X[P,N][3:0] & FCH ports: P_UMI_[T,R]X[P,N][3:0]
0121h	Gfx Links[7:0]: P_GFX_[T,R]X[P,N][7:0]
0221h	Gfx Links[15:8]: P_GFX_[T,R]X[P,N][15:8]
0122h	DDI 1: DP1_TX[P,N][3:0] & DDI 0: DP0_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:0]

Bits	Description
31:28	Reserved.
27:23	RttRawCal: receiver termination resistance (Rtt) raw calibration value. Read-only; Updated-by-hardware. Reset: 0. This field provides the raw Rtt calibration value as determined by the compensation circuit.
22:18	RonRawCal: transmitter resistance (Ron) raw calibration value. Read-only; Updated-by-hardware. Reset: 0. This field provides the raw Ron calibration value as determined by the compensation circuit.
17:0	Reserved.

D0F0xE4_x0[2:1]2[3:0]_000[2:1] Phy Impedance Control

Updates to these registers that result in a change to impedance may not take effect in the phy for up to 2 micro-

seconds after the update to this register completes.

Table 78: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000[2:1]

D0F0xE0[31:16]	D0F0xE0[15:0]	
	0002h	0001h
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0]	FCH ports: P_UMI_[T,R]X[P,N][3:0]
0121h	Graphics port lower: P_GFX_[T,R]X[P,N][7:4]	Graphics port lower: P_GFX_[T,R]X[P,N][3:0]
0221h	Graphics port upper: P_GFX_[T,R]X[P,N][15:12]	Graphics port upper: P_GFX_[T,R]X[P,N][11:8]
0122h	DDI 0: DP0_TX[P,N][3:0]	DDI 1: DP1_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:4]	DDI 2: DP2_TX[P,N][3:0]

Bits	Description														
31:29	<p>RttCtl: receiver termination resistance (Rtt) control. Read-write. Reset: 0. This field specifies how the receiver termination resistance value is calculated. All values between 00h and 1Fh are valid.</p> <table border="0"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Rtt is as determined by the compensation circuit, D0F0xE4_x0[2:1]2[3:0]_0000[RttRawCal].</td> </tr> <tr> <td>001b</td> <td>Rtt is as specified by (RttIndex - 3).</td> </tr> <tr> <td>010b</td> <td>Rtt is as specified by the difference: RttRawCal - RttIndex. If this results in a value that is less than 00h, then 00h is used.</td> </tr> <tr> <td>011b</td> <td>Rtt is as specified by the sum: RttRawCal + RttIndex. If this results in a value that is greater than 1Fh, then 1Fh is used.</td> </tr> <tr> <td>100b</td> <td>Enable only one tap of the Rtt resistor, as specified by RttIndex, and disable the base resistor that is normally always enabled. This is intended for testing purposes only.</td> </tr> <tr> <td>111b-101b</td> <td>Reserved.</td> </tr> </tbody> </table> <p>For all modes (except 100b), higher values reduce the resistance of Rtt and lower values increase the resistance of Rtt. If RttCtl is programmed to either 011b or 100b, the value of RttRawCal + RttIndex must be less than or equal to 24.</p>	<u>Bits</u>	<u>Definition</u>	000b	Rtt is as determined by the compensation circuit, D0F0xE4_x0[2:1]2[3:0]_0000[RttRawCal].	001b	Rtt is as specified by (RttIndex - 3).	010b	Rtt is as specified by the difference: RttRawCal - RttIndex. If this results in a value that is less than 00h, then 00h is used.	011b	Rtt is as specified by the sum: RttRawCal + RttIndex. If this results in a value that is greater than 1Fh, then 1Fh is used.	100b	Enable only one tap of the Rtt resistor, as specified by RttIndex, and disable the base resistor that is normally always enabled. This is intended for testing purposes only.	111b-101b	Reserved.
<u>Bits</u>	<u>Definition</u>														
000b	Rtt is as determined by the compensation circuit, D0F0xE4_x0[2:1]2[3:0]_0000[RttRawCal].														
001b	Rtt is as specified by (RttIndex - 3).														
010b	Rtt is as specified by the difference: RttRawCal - RttIndex. If this results in a value that is less than 00h, then 00h is used.														
011b	Rtt is as specified by the sum: RttRawCal + RttIndex. If this results in a value that is greater than 1Fh, then 1Fh is used.														
100b	Enable only one tap of the Rtt resistor, as specified by RttIndex, and disable the base resistor that is normally always enabled. This is intended for testing purposes only.														
111b-101b	Reserved.														
28:21	Reserved														
20:16	RttIndex: receiver termination resistance (Rtt) index. Read-write. Reset: 0. See RttCtl.														

15:13	<p>RonCtl: transmitter resistance (Ron) control. Read-write. Reset: 0. This field specifies how the transmitter resistance value is calculated.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Ron is as determined by the compensation circuit, D0F0xE4_x0[2:1]2[3:0]_0000[RonRawCal].</td> </tr> <tr> <td>001b</td> <td>Ron is as specified by the RonIndex field.</td> </tr> <tr> <td>010b</td> <td>Ron is as specified by the difference: RonRawCal - RonIndex. If this results in a value that is less than 00h, then 00h is used.</td> </tr> <tr> <td>011b</td> <td>Ron is as specified by the sum: RonRawCal + RonIndex. If this results in a value that is greater than 1Fh, then 1Fh is used.</td> </tr> <tr> <td>100b</td> <td>Enable only one tap of the Ron resistor, as specified by RonIndex, and disable the base resistor that is normally always enabled. This is intended for testing purposes only.</td> </tr> <tr> <td>111b-101b</td> <td>Reserved.</td> </tr> </tbody> </table> <p>For all modes (except 100b), higher values reduce the resistance of Ron and lower values increase the resistance of Ron. If RonCtl is programmed to either 011b or 100b, the value of RonRawCal + RonIndex must be less than or equal to 23.</p>	Bits	Definition	000b	Ron is as determined by the compensation circuit, D0F0xE4_x0[2:1]2[3:0]_0000[RonRawCal] .	001b	Ron is as specified by the RonIndex field.	010b	Ron is as specified by the difference: RonRawCal - RonIndex. If this results in a value that is less than 00h, then 00h is used.	011b	Ron is as specified by the sum: RonRawCal + RonIndex. If this results in a value that is greater than 1Fh, then 1Fh is used.	100b	Enable only one tap of the Ron resistor, as specified by RonIndex, and disable the base resistor that is normally always enabled. This is intended for testing purposes only.	111b-101b	Reserved.
Bits	Definition														
000b	Ron is as determined by the compensation circuit, D0F0xE4_x0[2:1]2[3:0]_0000[RonRawCal] .														
001b	Ron is as specified by the RonIndex field.														
010b	Ron is as specified by the difference: RonRawCal - RonIndex. If this results in a value that is less than 00h, then 00h is used.														
011b	Ron is as specified by the sum: RonRawCal + RonIndex. If this results in a value that is greater than 1Fh, then 1Fh is used.														
100b	Enable only one tap of the Ron resistor, as specified by RonIndex, and disable the base resistor that is normally always enabled. This is intended for testing purposes only.														
111b-101b	Reserved.														
12:5	Reserved.														
4:0	RonIndex: transmitter resistance (Ron) index. Read-write. Reset: 0. See RonCtl.														

D0F0xE4_x0[2:1]2[3:0]_000[A:9] Phy Clock Tree Control

BIOS: See 2.11.4.2.2 [Link Configuration and Core Initialization].

Table 79: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000[A:9]

D0F0xE0[31:16]	D0F0xE0[15:0]	
	000Ah	0009h
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0]	FCH ports: P_UMI_[T,R]X[P,N][3:0]
0121h	Graphics port upper: P_GFX_[T,R]X[P,N][15:12]	Graphics port upper: P_GFX_[T,R]X[P,N][11:8]
0221h	Graphics port lower: P_GFX_[T,R]X[P,N][7:4]	Graphics port lower: P_GFX_[T,R]X[P,N][3:0]
0122h	DDI 0: DP0_TX[P,N][3:0]	DDI 1: DP1_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:4]	DDI 2: DP2_TX[P,N][3:0]

Bits	Description
31	PCIEPIISel. Read-write. Reset: 1. 1=Selects PCIe clock. 0=Selects the display clock.
30:29	Reserved.
28	CascadedPIISel. Read-write. Reset: 0. 1=Selects cascaded PLL clocks. 0=Selects the external display PLL clocks.
27:26	Reserved.
25	DisplayStream. Read-write. Reset: 0. 1=Selects display PLL1 and cascaded PLL clock 1. 0=Selects display PLL0 and cascaded PLL clock 0.

24	ClkOff. Read-write. Reset: 0. 1=Turns off clock to 4 lane wide sub-link within phy. 0=Enables clock to 4 lane wide sublink within phy.
23:0	Reserved.

D0F0xE4_x0[2:1]2[3:0]_000[C:B] Phy Serial Bus Packet Control

This register provides control to enable or disable various fields contained in the phy serial bus primary control packet, the margining packet and the miscellaneous control packet.

Table 80: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000[C:B]

D0F0xE0[31:16]	D0F0xE0[15:0]	
	000Ch	000Bh
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0]	FCH ports: P_UMI_[T,R]X[P,N][3:0]
0121h	Graphics port lower: P_GFX_[T,R]X[P,N][7:4]	Graphics port lower: P_GFX_[T,R]X[P,N][3:0]
0221h	Graphics port upper: P_GFX_[T,R]X[P,N][15:12]	Graphics port upper: P_GFX_[T,R]X[P,N][11:8]
0122h	DDI 0: DP0_TX[P,N][3:0]	DDI 1: DP1_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:4]	DDI 2: DP2_TX[P,N][3:0]

Bits	Description
31:16	Reserved.
15	PllCmpPktSbiEn. Read-write. Reset: 1. 1=Enables the serial bus PLL component packet used for controlling PLL features such as mode of operation and power states.
14	MargPktSbiEn. Read-write. Reset: 1. IF (REG==D0F0xE4_x012[3:2]_000[C:B]) THEN BIOS: 0. ENDIF. 1=Enables the serial bus margining update packet used for controlling PCIe transmit margining test.
13:9	Reserved.
8	EiDetSbiEn. Read-write. Reset: 1. 1=Enables the electrical idle detector control field in the primary control packet.
7	IncoherentClkSbiEn. Read-write. Reset: 1. 1=Enables the incoherent clock control field in the primary control packet.
6	SkipBitSbiEn. Read-write. Reset: 1. 1=Enables the skip bit control field in the primary control packet.
5	OffsetCancelSbiEn. Read-write. Reset: 1. 1=Enables the offset cancellation control field in the primary control packet.
4	DllLockSbiEn. Read-write. Reset: 1. 1=Enables the DLL lock control field in the primary control packet.
3	FreqDivSbiEn. Read-write. Reset: 1. 1=Enables the frequency divider control field in the primary control packet.
2	PcieModeSbiEn. Read-write. Reset: 1. IF (REG==D0F0xE4_x012[3:2]_000[C:B]) THEN BIOS: 0. ENDIF 1=Enables the phy mode control field in the primary control packet.
1	RxPwrSbiEn. Read-write. Reset: 1. 1=Enables the Rx power state control field in the primary control packet.
0	TxPwrSbiEn. Read-write. Reset: 1. 1=Enables the Tx power state control field in the primary control packet.

D0F0xE4_x0[2:1]2[3:0]_000D Phy Serial Bus Compensation Component Packet Enable

This register provides control to enable or disable the phy serial bus compensation component packet.

Table 81: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_000D

D0F0xE0[31:16]	D0F0xE0[15:0]
	000Dh
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0] & FCH ports: P_UMI_[T,R]X[P,N][3:0]
0121h	Gfx Links[7:0]: P_GFX_[T,R]X[P,N][7:0]
0221h	Gfx Links[15:8]: P_GFX_[T,R]X[P,N][15:8]
0122h	DDI 1: DP1_TX[P,N][3:0] & DDI 0: DP0_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:0]

Bits	Description
31:1	Reserved.
0	CmpCmpPktSbiEn. Read-write. Reset: 1. 1=Enables the serial bus compensation component packet used for controlling compensation circuit features and power states.

D0F0xE4_x0[2:1]2[3:0]_2000 Phy PLL Power State Control

This register provides control of the phy PLL component power state.

Table 82: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_2000

D0F0xE0[31:16]	D0F0xE0[15:0]
	2000h
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0] & FCH port: P_UMI_[T,R]X[P,N][3:0]
0121h	Gfx Links[7:0]: P_GFX_[T,R]X[P,N][7:0]
0221h	Gfx Links[15:8]: P_GFX_[T,R]X[P,N][15:8]
0122h	DDI 1: DP1_TX[P,N][3:0] & DDI 0: DP0_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:0]

Bits	Description												
31:4	Reserved.												
3	PlIAutoPwrDownDis. Read-write. Reset: 0. 1=Disables the automatic power down feature. 0=Enables the automatic power down feature; PLL powers down when it determines that it is unused.												
2:0	<p>PlIPowerDownEn. Read-write. Reset: 0. This field specifies the power state of the phy PLL circuit component.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>L0 power state; all circuits are enabled.</td> </tr> <tr> <td>001b</td> <td>LS1 power state; all circuits are enabled.</td> </tr> <tr> <td>010b</td> <td>LS2 power state; PLL is powered down and clock tree is gated off.</td> </tr> <tr> <td>110b-011b</td> <td>Reserved.</td> </tr> <tr> <td>111b</td> <td>PHYOFF power state; all circuits are disabled to achieve the lowest power consumption.</td> </tr> </tbody> </table>	Bits	Definition	000b	L0 power state; all circuits are enabled.	001b	LS1 power state; all circuits are enabled.	010b	LS2 power state; PLL is powered down and clock tree is gated off.	110b-011b	Reserved.	111b	PHYOFF power state; all circuits are disabled to achieve the lowest power consumption.
Bits	Definition												
000b	L0 power state; all circuits are enabled.												
001b	LS1 power state; all circuits are enabled.												
010b	LS2 power state; PLL is powered down and clock tree is gated off.												
110b-011b	Reserved.												
111b	PHYOFF power state; all circuits are disabled to achieve the lowest power consumption.												

D0F0xE4_x0[2:1]2[3:0]_2002 Phy PLL Control

The newly written values to this register do not take effect until an explicit shadow update event takes place via setting [D0F0xE4_x0\[2:1\]2\[3:0\]_2008\[PIIControlUpdate\]](#), or until the PLL component exits the LS2 or PHYOFF power state. Any read from this register always returns the current register value, not the value pending until the next update event. See [D0F0xE4_x0\[2:1\]2\[3:0\]_2008\[PIIControlUpdate\]](#).

Table 83: Index Mapping for [D0F0xE4_x0\[2:1\]2\[3:0\]_2002](#)

D0F0xE0[31:16]	D0F0xE0[15:0]
	2002h
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0] & FCH port: P_UMI_[T,R]X[P,N][3:0]
0121h	Graphics port upper: P_GFX_[T,R]X[P,N][15:8]
0221h	Graphics port lower: P_GFX_[T,R]X[P,N][7:0]
0122h	DDI 0: DP0_TX[P,N][3:0] & DDI 1: DP1_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:0]

Bits	Description
31	IsLc: PLL select. Read-write. Reset: 0. 1=Selects the LC tank raw PLL. 0=Selects the ring oscillator raw PLL.
30:28	Reserved.
27	RoCalEn: Ring oscillator calibration enable. Read-write. Reset: 0. BIOS: IF (REG==D0F0xE4_x0120_2002) THEN 1 ELSE 0 ENDIF. A 0 to 1 edge transition of this bit triggers a calibration cycle for the ring oscillator VCO. Reset and set this bit again to trigger another calibration cycle if needed.
26:0	Reserved.

D0F0xE4_x0[2:1]2[3:0]_2005 Phy PLL Frequency and Mode Control

This register provides PLL operation mode control and frequency selection. The newly written values to this register do not take effect until an explicit shadow update event takes place via setting [D0F0xE4_x0\[2:1\]2\[3:0\]_2008\[PIIControlUpdate\]](#), or until the PLL component exits the LS2 or PHYOFF power state. Any read from this register always returns the current register value, not the value pending until the next update event. See [D0F0xE4_x0\[2:1\]2\[3:0\]_2008\[PIIControlUpdate\]](#).

Table 84: Index Mapping for [D0F0xE4_x0\[2:1\]2\[3:0\]_2005](#)

D0F0xE0[31:16]	D0F0xE0[15:0]
	2005h
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0] & FCH port: P_UMI_[T,R]X[P,N][3:0]
0121h	Gfx Links[7:0]: P_GFX_[T,R]X[P,N][7:0]
0221h	Gfx Links[15:8]: P_GFX_[T,R]X[P,N][15:8]
0122h	DDI 1: DP1_TX[P,N][3:0] & DDI 0: DP0_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:0]

Bits	Description																												
31:15	Reserved.																												
14:13	PllMode: PLL operation mode. Read-write. Reset: 0. BIOS: See 2.11.4.2.1 [Clock Configuration]. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>PCIe mode</td> </tr> <tr> <td>01b</td> <td>Disabled</td> </tr> <tr> <td>10b</td> <td>IsCascadedClk0 mode</td> </tr> <tr> <td>11b</td> <td>IsCascadedClk1 mode</td> </tr> </tbody> </table>	Bits	Definition	00b	PCIe mode	01b	Disabled	10b	IsCascadedClk0 mode	11b	IsCascadedClk1 mode																		
Bits	Definition																												
00b	PCIe mode																												
01b	Disabled																												
10b	IsCascadedClk0 mode																												
11b	IsCascadedClk1 mode																												
12:11	Reserved.																												
10:9	PllClkFreqExt. Read-write. Reset: 10b. See PllClkFreq.																												
8:4	Reserved.																												
3:0	PllClkFreq: PLL frequency select. Read-write. Reset: 0011b. BIOS: See 2.11.4.2.1 [Clock Configuration]. This field is used together with PllClkFreqExt to specify the PLL frequency. <table border="1"> <thead> <tr> <th>{ PllClkFreqExt, PllClkFreq }</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>100000b</td> <td>810MHz (display port)</td> </tr> <tr> <td>100001b</td> <td>2.7GHz (display port)</td> </tr> <tr> <td>100010b</td> <td>1.35GHz (display port)</td> </tr> <tr> <td>100011b</td> <td>2.5GHz (PCIe Gen 2)</td> </tr> <tr> <td>110000b</td> <td>125MHz - 156.25MHz (DVI/HDMI)</td> </tr> <tr> <td>110001b</td> <td>156.25MHz - 206.25MHz (DVI/HDMI)</td> </tr> <tr> <td>110010b</td> <td>206.25MHz - 250MHz (DVI/HDMI)</td> </tr> <tr> <td>110011b</td> <td>250MHz - 312.5MHz (DVI/HDMI)</td> </tr> <tr> <td>110100b</td> <td>312.5MHz - 412.5MHz (DVI/HDMI)</td> </tr> <tr> <td>110101b</td> <td>412.5MHz - 500MHz (DVI/HDMI)</td> </tr> <tr> <td>110110b</td> <td>500MHz - 625MHz (DVI/HDMI)</td> </tr> <tr> <td>110111b</td> <td>625MHz - 825MHz (DVI/HDMI)</td> </tr> <tr> <td>all others</td> <td>Reserved</td> </tr> </tbody> </table>	{ PllClkFreqExt, PllClkFreq }	Definition	100000b	810MHz (display port)	100001b	2.7GHz (display port)	100010b	1.35GHz (display port)	100011b	2.5GHz (PCIe Gen 2)	110000b	125MHz - 156.25MHz (DVI/HDMI)	110001b	156.25MHz - 206.25MHz (DVI/HDMI)	110010b	206.25MHz - 250MHz (DVI/HDMI)	110011b	250MHz - 312.5MHz (DVI/HDMI)	110100b	312.5MHz - 412.5MHz (DVI/HDMI)	110101b	412.5MHz - 500MHz (DVI/HDMI)	110110b	500MHz - 625MHz (DVI/HDMI)	110111b	625MHz - 825MHz (DVI/HDMI)	all others	Reserved
{ PllClkFreqExt, PllClkFreq }	Definition																												
100000b	810MHz (display port)																												
100001b	2.7GHz (display port)																												
100010b	1.35GHz (display port)																												
100011b	2.5GHz (PCIe Gen 2)																												
110000b	125MHz - 156.25MHz (DVI/HDMI)																												
110001b	156.25MHz - 206.25MHz (DVI/HDMI)																												
110010b	206.25MHz - 250MHz (DVI/HDMI)																												
110011b	250MHz - 312.5MHz (DVI/HDMI)																												
110100b	312.5MHz - 412.5MHz (DVI/HDMI)																												
110101b	412.5MHz - 500MHz (DVI/HDMI)																												
110110b	500MHz - 625MHz (DVI/HDMI)																												
110111b	625MHz - 825MHz (DVI/HDMI)																												
all others	Reserved																												

D0F0xE4_x0[2:1]2[3:0]_2008 Phy PLL Update Control

Table 85: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_2008

D0F0xE0[31:16]	D0F0xE0[15:0]
	2008h
0120h	GPP ports: P_GPP_[T,R]X[P,N][3:0] & FCH port: P_UMI_[T,R]X[P,N][3:0]
0121h	Gfx Links[7:0]: P_GFX_[T,R]X[P,N][7:0]
0221h	Gfx Links[15:8]: P_GFX_[T,R]X[P,N][15:8]
0122h	DDI 1: DP1_TX[P,N][3:0] & DDI 0: DP0_TX[P,N][3:0]
0123h	DDI 2: DP2_TX[P,N][6:0]

Bits	Description
31:30	Reserved.

29	VdDetectEn: PLL output clock gating enable. Read-write. Reset: 0. BIOS: 1. 1=Enables clock gating during LS2 exit. The internally regulated power supply to the phy is turned off and then back on during LS2 entry and exit. A power sniffer circuit detects the internally regulated power supply level. PLL output clock distribution can be gated off before the regulated power supply reaches the desired level during power up, this effectively prevents possible reliability issue arising from undesirable drivers contention and crow bar current stress during power up upon LS2 exit. Clock distribution is turned on only when the regulated power reaches the right level. 0=PLL output clock distribution is not gated during power up upon LS2 exit.																		
28:26	Reserved.																		
25:23	MeasCycCntVal[2:0]. Read-write. Reset: 0. This field specifies the time interval over which the LC VCO feedback clock cycles is counted. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>3167 count for 31.5kHz modulation</td> </tr> <tr> <td>001b</td> <td>3023 count for 33.0kHz modulation (max spec)</td> </tr> <tr> <td>010b</td> <td>3325 count for 33.0 kHz modulation (min spec)</td> </tr> <tr> <td>011b</td> <td>6334 count, 2 times default (long count)</td> </tr> <tr> <td>100b</td> <td>12667 count, 4 times default (very long count)</td> </tr> <tr> <td>101b</td> <td>2048, 11bits (shorter count for debug)</td> </tr> <tr> <td>110b</td> <td>1024, 10bits</td> </tr> <tr> <td>111b</td> <td>512, 9bits</td> </tr> </tbody> </table>	Bits	Definition	000b	3167 count for 31.5kHz modulation	001b	3023 count for 33.0kHz modulation (max spec)	010b	3325 count for 33.0 kHz modulation (min spec)	011b	6334 count, 2 times default (long count)	100b	12667 count, 4 times default (very long count)	101b	2048, 11bits (shorter count for debug)	110b	1024, 10bits	111b	512, 9bits
Bits	Definition																		
000b	3167 count for 31.5kHz modulation																		
001b	3023 count for 33.0kHz modulation (max spec)																		
010b	3325 count for 33.0 kHz modulation (min spec)																		
011b	6334 count, 2 times default (long count)																		
100b	12667 count, 4 times default (very long count)																		
101b	2048, 11bits (shorter count for debug)																		
110b	1024, 10bits																		
111b	512, 9bits																		
22:1	Reserved.																		
0	PlIControlUpdate: PLL control register update. Read-write. Reset: 0. 0 to 1 transition of this bit triggers the explicit shadow update event for the phy PLL registers that are shadowed.																		

3.3.3.2 Phy Receiver Lane Control Registers

Each receiver lane has a group of registers for controlling the operation of the lane. The mapping from address register to receiver lane is shown in [Table 86](#). Multiple receiver lanes may be written at the same time using per nibble and per byte broadcast write addresses. The mapping from broadcast address to receiver lanes is shown in [Table 87](#).

Table 86: Phy per receiver lane register addresses

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		438xh	430xh	428xh	420xh	418xh	410xh	408xh	400xh
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 87: Phy receiver broadcast register addresses

D0F0xE0[31:16]	D0F0xE0[15:0]		
	57[1,0]xh	56[1,0]xh	50[1,0]xh
0120h	P_GPP_RX[P,N][3:0]	P_UMI_RX[P,N][3:0]	P_GPP_RX[P,N][3:0], P_UMI_RX[P,N][3:0]
0121h	P_GFX_RX[P,N][7:4]	P_GFX_RX[P,N][3:0]	P_GFX_RX[P,N][7:0]
0221h	P_GFX_RX[P,N][15:12]	P_GFX_RX[P,N][11:8]	P_GFX_RX[P,N][15:8]

D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]1 Phy Receiver DLL Control and Test 1

These registers provide control of the DLL and duty cycle correction circuit associated with the receive lanes on the phys.

Table 88: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]1

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		4381h	4301h	4281h	4201h	4181h	4101h	4081h	4001h
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 89: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]1

D0F0xE0[31:16]	D0F0xE0[15:0]		
	5701h	5601h	5001h
0120h	D0F0xE4_x0120_4[3:2][8,0]1	D0F0xE4_x0120_4[1:0][8,0]1	D0F0xE4_x0120_4[3:0][8,0]1
0121h	D0F0xE4_x0121_4[3:2][8,0]1	D0F0xE4_x0121_4[1:0][8,0]1	D0F0xE4_x0121_4[3:0][8,0]1
0221h	D0F0xE4_x0221_4[3:2][8,0]1	D0F0xE4_x0221_4[1:0][8,0]1	D0F0xE4_x0221_4[3:0][8,0]1

Bits	Description
31:16	Reserved.
15	ForceDccRecalc: Force DCC code recalculation. Read-write. Reset: 0. BIOS: IF (REG==D0F0xE4_x0120_4[3:0][8,0]1) THEN 1 ELSE 0 ENDIF. A 0 to 1 edge transition of this bit forces the DCC code to be recalculated when the DLL core loop is locked the next time. This could be useful for debug. When transitioning between the RO PLL to the LC PLL this bit should be set.
14:0	Reserved.

D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]2 Phy Receiver Phase Loop Filter Control




Figure 10: Phy recovered clock and sample clock

When the link is in a mode that relies on dynamic phase alignment (automatic sample-clock correction), then the processor generates a recovered clock for each lane based on transitions in the lane. The ideal recovered clock transitions at exactly the same time as the transitions in the lane. Phase detection logic detects if the recovered clock transitions before or after the lane transition. The digital loop filter (DLF) is logic that adjusts

the phase of the recovered clock such that its transitions match the transition time of the lane as much as possible. The DLF counts the number of times the lane transitions before the recovered clock versus after to determine whether the recovered clock phase requires adjustment. The DLF uses an 8-bit counter, called the loop filter counter (LFC) for this purpose. The LFC controls are included in this register. They specify DLF behavior as follows:

- LfcMax is programmed to be greater than LfcMin.
- The LFC is initialized to LfcMin.
- The LFC is updated periodically. The logic keeps a tally of the number of lane transitions occurring before and after the recovered clock transition within each update period.
- To start, if there is a net lane transition occurs after the recovered clock transition within the update period, the LFC is incremented by the net value; on the other hand, if there is a net lane transition occurs before the recovered clock transition, the LFC is decremented. However, if the LFC is ever decremented while it is zero, these rules are reversed (and the LFC is incremented instead). Thus, if there is a phase correction needed, the LFC trends either upward or downward; if it trends downward, it hits zero and then trends upward again.
- If the LFC reaches LfcMax value, then (1) the phase of the recovered clock is adjusted in the appropriate direction, (2) the LFC is set to the LfcMin value.

The LfcMin and LfcMax fields are designed to improve the stability of the recovered clock phase while improving the response time for multiple phase updates in the same direction. For example, if the recovered clock phase needs several adjustments in the same direction, then the LFC increments until it hits LfcMax value and then be set to LfcMin (and trigger a phase adjustment); then it would increment to LfcMax value again to trigger the next phase adjustment. If, however, the next phase adjustment needs to be in the opposite direction, the LFC would decrement to zero, change direction, and then increment up to LfcMax again. In this way, phase adjustments in the same direction occur more quickly than phase adjustments in the opposite direction of the prior phase adjustment.

The nominal sample clock is offset by 90 degrees from the recovered clock. An offset can be inserted to move the sample clock from the nominal position, based on SamClkPiOffset and SamClkPiOffsetSign.

Table 90: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]2

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		4382h	4302h	4282h	4202h	4182h	4102h	4082h	4002h
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 91: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]2

D0F0xE0[31:16]	D0F0xE0[15:0]		
	5702h	5602h	5002h
0120h	D0F0xE4_x0120_4[3:2][8,0]2	D0F0xE4_x0120_4[1:0][8,0]2	D0F0xE4_x0120_4[3:0][8,0]2
0121h	D0F0xE4_x0121_4[3:2][8,0]2	D0F0xE4_x0121_4[1:0][8,0]2	D0F0xE4_x0121_4[3:0][8,0]2
0221h	D0F0xE4_x0221_4[3:2][8,0]2	D0F0xE4_x0221_4[1:0][8,0]2	D0F0xE4_x0221_4[3:0][8,0]2

Bits	Description
31:30	Reserved.
29:22	LfcMax: loop filter counter maximum value. Read-write. Reset: 08h. BIOS: 08h.
21:14	LfcMin: loop filter counter minimum value. Read-write. Reset: 00h. BIOS: 00h.
13:8	Reserved.
7	SamClkPiOffsetEn: sample clock phase interpolator offset enable. Read-write. Reset: 0. 1=Enable offset insertion around the nominal sample clock position.
6:4	SamClkPiOffset: sample clock phase interpolator offset setting. Read-write. Reset: X. This field specifies the magnitude of the offset of the sample clock from the nominal position. See Figure 10 . This field is encoded as follows. <ul style="list-style-type: none"> • Sample clock phase interpolator offset = (SamClkPiOffset + 1) * step size. • If link speed is >3.6GT/s, the expected typical step size is 2ps with a +/-1ps error. • If link speed is <=3.6GT/s, the expected typical step size is 3ps with a +/-1ps error.
3	SamClkPiOffsetSign: sample clock phase interpolator offset setting sign bit. Read-write. Reset: X. 0=Sample clock is moved to before the nominal position. 1=Sample clock is moved to after the nominal position. See SamClkPiOffset and Figure 10 .
2:0	Reserved.

D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]5 Phy Receiver Timing Margin Test

The built in jitter injection test mode is useful for checking the clock data recovery tracking bandwidth of the receiver. By forcing the sample clock to move from the lock position by a controlled amount and then observing the time it takes to recover, the tracking rate and bandwidth can be estimated. This register provides the control of the test mode.

The jitter injection test mode works as follows.

- The circuit is clocked by a jitter injection clock derived from dividing the link forwarded clock by 2.5; for example, if the link speed is 5.2GT/s and the link forwarded clock frequency is 2.6GHz, the jitter injection clock frequency becomes 1.04GHz.
- There are 2 phases, the on phase and the off phase. It starts with the on phase once the test mode is enabled.
- During the on phase, at every tick of jitter injection clock, the sample clock is moved away from the nominal lock position by $1/96 * UI$.
- The direction of adjustment is specified by JitterInjDir.
- The on phase adjustment continues for a number of times as specified by JitterInjOnCnt.
- Then the adjustment turns off for a duration specified by {JitterInjOffCnt, JitterInjOnCnt} * jitter injection clock period, this is known as the off phase. During this time, clock data recovery resumes to try to adjust the position of the sample clock back to the center of the data eye.
- The off phase is followed by the on phase again. The process continues to alternate between the on phase and the off phase until the jitter injection test mode is disabled.

In addition, the JitterInjHold bit may be set to inject a hold state at the end of the on phase. This stops clock data recovery from resuming after the on phase, hence holding the sample clock at its last adjusted position until the JitterInjHold bit is cleared. This test mode may be useful for margining the width of the input data eye.

This margining mechanism is not characterized for precision jitter adjustments or measurements.

Table 92: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]5

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		4385h	4305h	4285h	4205h	4185h	4105h	4085h	4005h
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 93: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]5

D0F0xE0[31:16]	D0F0xE0[15:0]		
	5705h	5605h	5005h
0120h	D0F0xE4_x0120_4[3:2][8,0]5	D0F0xE4_x0120_4[1:0][8,0]5	D0F0xE4_x0120_4[3:0][8,0]5
0121h	D0F0xE4_x0121_4[3:2][8,0]5	D0F0xE4_x0121_4[1:0][8,0]5	D0F0xE4_x0121_4[3:0][8,0]5
0221h	D0F0xE4_x0221_4[3:2][8,0]5	D0F0xE4_x0221_4[1:0][8,0]5	D0F0xE4_x0221_4[3:0][8,0]5

Bits	Description						
31	Reserved.						
30	JitterInjEn: jitter injection enable. Read-write. Reset: 0. 1=Jitter injection test mode is enabled.						
29	JitterInjDir: jitter injection direction. Read-write. Reset: 0. <table border="0"> <tr> <td><u>Bit</u></td> <td><u>Definition</u></td> </tr> <tr> <td>0</td> <td>Move clock before the nominal lock position.</td> </tr> <tr> <td>1</td> <td>Move clock after the nominal lock position.</td> </tr> </table>	<u>Bit</u>	<u>Definition</u>	0	Move clock before the nominal lock position.	1	Move clock after the nominal lock position.
<u>Bit</u>	<u>Definition</u>						
0	Move clock before the nominal lock position.						
1	Move clock after the nominal lock position.						
28:23	JitterInjOnCnt: jitter injection on count. Read-write. Reset: 0.						
22:16	Reserved.						
15:10	JitterInjOffCnt: jitter injection off count. Read-write. Reset: 0. The jitter injection off time count is a 12bit code, this field specifies the most significant 6 bits. The least significant 6 bits are the same as JitterInjOnCnt.						
9	JitterInjHold: jitter injection hold. Read-write. Reset: 0. 1=Jitter injection hold is enabled.						
8:0	Reserved.						

D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]6 Phy Receiver DFE and DFR Control

The processor supports decision feedback restore (DFR), a function that enables on-chip AC coupling on the receiver path, to improve the receiver's ability to operate over a longer channel. In this mode, the receiver on the processor must be programmed with the expected peak single-ended DC voltage level over the single-ended DC common mode voltage level, as seen by the receiver, when a static 1 or 0 is driven. For example, without deemphasis at nominal supply voltage of 1.2V, the peak single ended voltage is expected to be 300mV ideally above the single ended DC common mode voltage level. The value is dependent on the deemphasis setting of the transmitter on the other end of the channel.

Table 94: Recommended DCV settings

Far-device deemphasis setting	DCV
-------------------------------	-----

Table 94: Recommended DCV settings

No deemphasis	20h
-3dB postcursor	17h
-6dB postcursor	10h

Decision feedback equalization (DFE) can be enabled to enhance link operation. Once enabled, the receiver uses the logic level of the previous data bit to adjust the voltage threshold of the sampler in the direction that causes the sampler to switch sooner when the data bit transitions to the opposite logic level for the next bit. The control and DFE voltage level are included in this register.

Table 95: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]6

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		4386h	4306h	4286h	4206h	4186h	4106h	4086h	4006h
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 96: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]6

D0F0xE0[31:16]	D0F0xE0[15:0]		
	5706h	5606h	5006h
0120h	D0F0xE4_x0120_4[3:2][8,0]6	D0F0xE4_x0120_4[1:0][8,0]6	D0F0xE4_x0120_4[3:0][8,0]6
0121h	D0F0xE4_x0121_4[3:2][8,0]6	D0F0xE4_x0121_4[1:0][8,0]6	D0F0xE4_x0121_4[3:0][8,0]6
0221h	D0F0xE4_x0221_4[3:2][8,0]6	D0F0xE4_x0221_4[1:0][8,0]6	D0F0xE4_x0221_4[3:0][8,0]6

Bits	Description												
31:8	Reserved.												
7	DfeEn: DFE enable. Read-write. Reset: 0. 1=Decision feedback equalization is enabled.												
6:5	DfeVoltage: DFE offset voltage level. Read-write. Reset: 0. This field specifies the magnitude of the DFE offset voltage. <table border="0"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>DFE offset voltage=25mV.</td> <td>10b</td> <td>DFE offset voltage=12.5mV.</td> </tr> <tr> <td>01b</td> <td>DFE offset voltage=0mV.</td> <td>11b</td> <td>DFE offset voltage=31.25mV.</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	00b	DFE offset voltage=25mV.	10b	DFE offset voltage=12.5mV.	01b	DFE offset voltage=0mV.	11b	DFE offset voltage=31.25mV.
Bits	Definition	Bits	Definition										
00b	DFE offset voltage=25mV.	10b	DFE offset voltage=12.5mV.										
01b	DFE offset voltage=0mV.	11b	DFE offset voltage=31.25mV.										
4:0	Reserved.												

D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]A Phy DLL Test and Control 3**Table 97: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]A**

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		438Ah	430Ah	428Ah	420Ah	418Ah	410Ah	408Ah	400Ah
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0

Table 97: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]A

P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 98: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]A

D0F0xE0[31:16]	D0F0xE0[15:0]		
	570Ah	560Ah	500Ah
0120h	D0F0xE4_x0120_4[3:2][8,0]A	D0F0xE4_x0120_4[1:0][8,0]A	D0F0xE4_x0120_4[3:0][8,0]A
0121h	D0F0xE4_x0121_4[3:2][8,0]A	D0F0xE4_x0121_4[1:0][8,0]A	D0F0xE4_x0121_4[3:0][8,0]A
0221h	D0F0xE4_x0221_4[3:2][8,0]A	D0F0xE4_x0221_4[1:0][8,0]A	D0F0xE4_x0221_4[3:0][8,0]A

Bits	Description																
31:29	<p>Ls2ExitTime: LS2 exit time. Read-write. Reset: 0. BIOS: 001b. This field selects the internal timer that delays the turn-on of the DLL after exit from LS2 state to L0 state. The added delay allows the forwarded input clock to achieve better stability.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Delay=10us.</td> </tr> <tr> <td>001b</td> <td>Delay=5us.</td> </tr> <tr> <td>010b</td> <td>Delay=2.5us.</td> </tr> <tr> <td>011b</td> <td>Delay=1.25us.</td> </tr> <tr> <td>100b</td> <td>Delay=625ns.</td> </tr> <tr> <td>101b</td> <td>Delay=0s.</td> </tr> <tr> <td>11xb</td> <td>Reserved.</td> </tr> </tbody> </table> <p>The value specified by Ls2ExitTime must be less than the value specified by T0Time, or it can cause undefined behavior.</p>	Bits	Definition	000b	Delay=10us.	001b	Delay=5us.	010b	Delay=2.5us.	011b	Delay=1.25us.	100b	Delay=625ns.	101b	Delay=0s.	11xb	Reserved.
Bits	Definition																
000b	Delay=10us.																
001b	Delay=5us.																
010b	Delay=2.5us.																
011b	Delay=1.25us.																
100b	Delay=625ns.																
101b	Delay=0s.																
11xb	Reserved.																
28:18	Reserved.																
17	<p>DllLockFastModeEn: DLL lock fast mode enable. Read-write. Reset: 0. 1=Enables DLL lock fast mode. 0=DLL lock operates at standard speed.</p>																
16:15	Reserved.																
14:13	<p>AnalogWaitTime: analog wait time to turn on DLL. Read-write. Reset: 0. The turning on of the DLL circuit after cold reset is delayed by a timer specified by this field. The encodings are as follows:</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Delay=1.25us.</td> <td>10b</td> <td>Delay=2.5us.</td> </tr> <tr> <td>01b</td> <td>Delay=0.625us.</td> <td>11b</td> <td>Delay=0.3125us.</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	00b	Delay=1.25us.	10b	Delay=2.5us.	01b	Delay=0.625us.	11b	Delay=0.3125us.				
Bits	Definition	Bits	Definition														
00b	Delay=1.25us.	10b	Delay=2.5us.														
01b	Delay=0.625us.	11b	Delay=0.3125us.														
12:8	Reserved.																

7	BiasDisInLs2: bias disable in LS2 power state. Read-write. Reset: 0. IF ((REG==D0F0xE4_x0121_4[1:0][8,0]A && D0F0xE4_x0131_8040[OwnSlice]==0) (REG==D0F0xE4_x0121_4[3:2][8,0]A && D0F0xE4_x0131_8041[OwnSlice]==0) (REG==D0F0xE4_x0221_4[1:0][8,0]A && D0F0xE4_x0131_8042[OwnSlice]==0) (REG==D0F0xE4_x0221_4[3:2][8,0]A && D0F0xE4_x0131_8043[OwnSlice]==0) REG==D0F0xE4_x0120_4[1:0][8,0]A REG==D0F0xE4_x0120_4[3:2][8,0]A) THEN BIOS: 1. ENDIF. 1=Enables lower power LS2 state; current consumption is lowered by approximately 2.5mA per receive lane when compared to standard LS2 power mode. Setting this bit increases the amount of T0Time needed to relock the DLL. When this bit is set, Ls2ExitTime must be programmed to select a value that is greater than or equal to AnalogWaitTime. 0=Standard LS2 power mode.
6:5	Reserved.
4	LockDetOnLs2Exit: DLL lock detect on LS2 exit. Read-write. Reset: 0. This field selects the LS2 to L0 power state transition speed. 1=Fast transition mode selected. 0=Slow transition mode selected.
3:2	Reserved.
1	RxPcieMode: Receiver PCI Express Mode. Read-write. Reset: 0. Indicates whether the receiver lane is in Gen1 or Gen2 mode. 0=Gen1 mode. 1=Gen2 mode.
0	EnCoreLoopFirst: enable DLL core loop first on LS2 exit. Read-write. Reset: 0. This field selects LS2 to L0 power state transition speed. 1=Fast transition mode selected. 0=Slow transition mode selected.

D0F0xE4_x0[2:1]2[1:0]_ [5:4][7:6,3:0][8,0]F Phy Receiver DLL Test and Debug 5

Table 99: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_ [5:4][7:6,3:0][8,0]F

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		438Fh	430Fh	428Fh	420Fh	418Fh	410Fh	408Fh	400Fh
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 100: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_ [5:4][7:6,3:0][8,0]F

D0F0xE0[31:16]	D0F0xE0[15:0]		
	570Fh	560Fh	500Fh
0120h	D0F0xE4_x0120_4[3:2][8,0]F	D0F0xE4_x0120_4[1:0][8,0]F	D0F0xE4_x0120_4[3:0][8,0]F
0121h	D0F0xE4_x0121_4[3:2][8,0]F	D0F0xE4_x0121_4[1:0][8,0]F	D0F0xE4_x0121_4[3:0][8,0]F
0221h	D0F0xE4_x0221_4[3:2][8,0]F	D0F0xE4_x0221_4[1:0][8,0]F	D0F0xE4_x0221_4[3:0][8,0]F

Bits	Description
31:13	Reserved
12	DllProcessFreqCtlOverride. Read-write. Reset: 0. 1=Enables the override of DLL delay line capacitance settings with the values of DllProcessFreqCtlIndex1 and DllProcessFreqCtlIndex2.
11	Reserved.

10:7	DllProcessFreqCtlIndex2 . Read-write. Reset: 0. This field specifies the DLL delay line capacitance setting in PCIe Gen2 mode.
6:4	Reserved.
3:0	DllProcessFreqCtlIndex1 . Read-write. Reset: 0. This field specifies the DLL delay line capacitance setting in PCIe Gen1 mode.

3.3.3.3 Phy Transmitter Lane Control Registers

Each transmitter lane has a group of registers for controlling the operation of the lane. The mapping from address register to transmitter lane is shown in [Table 101](#). Multiple transmitter lanes may be written at the same time using per nibble and per byte broadcast write addresses. The mapping from broadcast address to transmitter lanes is shown in [Table 102](#).

Table 101: Phy per transmitter lane register addresses

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		638xh	630xh	628xh	620xh	618xh	610xh	608xh	600xh
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8
P_DP0_	0122h	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0	-	-	-	-
P_DP1_	0122h	-	-	-	-	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_DP2_	0123h	-	TX[P,N]6	TX[P,N]5	TX[P,N]4	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0

Table 102: Phy transmitter broadcast register addresses

D0F0xE0[31:16]	D0F0xE0[15:0]		
	77[1,0]xh	76[1,0]xh	70[1,0]xh
0120h	P_GPP_RX[P,N][3:0]	P_UMI_RX[P,N][3:0]	P_GPP_RX[P,N][3:0], P_UMI_RX[P,N][3:0]
0121h	P_GFX_RX[P,N][7:4]	P_GFX_RX[P,N][3:0]	P_GFX_RX[P,N][7:0]
0221h	P_GFX_RX[P,N][15:12]	P_GFX_RX[P,N][11:8]	P_GFX_RX[P,N][15:8]

D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]0 Phy Tx Deemphasis and Margining Control

Table 103: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]0

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		6380h	6300h	6280h	6200h	6180h	6100h	6080h	6000h
P_GPP_	0120h	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_GFX_	0121h	TX[P,N]7	TX[P,N]6	TX[P,N]5	TX[P,N]4	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_GFX_	0221h	TX[P,N]15	TX[P,N]14	TX[P,N]13	TX[P,N]12	TX[P,N]11	TX[P,N]10	TX[P,N]9	TX[P,N]8
P_DP0_	0122h	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0	-	-	-	-
P_DP1_	0122h	-	-	-	-	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_DP2_	0123h	-	TX[P,N]6	TX[P,N]5	TX[P,N]4	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0

Table 104: Broadcast Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]0

D0F0xE0[31:16]	D0F0xE0[15:0]		
	7700h	7600h	7000h
0120h	D0F0xE4_x0120_6[3:2][8,0]0	D0F0xE4_x0120_6[1:0][8,0]0	D0F0xE4_x0120_6[3:0][8,0]0
0121h	D0F0xE4_x0121_6[3:2][8,0]0	D0F0xE4_x0121_6[1:0][8,0]0	D0F0xE4_x0121_6[3:0][8,0]0
0221h	D0F0xE4_x0221_6[3:2][8,0]0	D0F0xE4_x0221_6[1:0][8,0]0	D0F0xE4_x0221_6[3:0][8,0]0
0122h	D0F0xE4_x0122_6[3:2][8,0]0	D0F0xE4_x0122_6[1:0][8,0]0	D0F0xE4_x0122_6[3:0][8,0]0
0123h	D0F0xE4_x0123_6[3:2][8,0]0	D0F0xE4_x0123_6[1:0][8,0]0	D0F0xE4_x0123_6[3:0][8,0]0

Bits	Description
31:16	Reserved.
15	DisLoImpIdle: disable low impedance idle. Read-write. Reset: 0. 1= Disables the low impedance electrical idle feature that requires both the true and complement pins of the transmitter to be pulled to VDDP/2 via low impedance termination in the range of 25 to 50 ohm upon entering electrical idle state. Instead, 5k ohm termination is used. 0=Enables low impedance electrical idle mode.
14:8	Reserved.
7	TxLs23ClkGateEn: LS2/LS3 clock gating enable. Read-write. Reset: 1. 1= Internal phy clock grids are gated during LS2 or PHY OFF states to save power.
6:4	Reserved.
3	Post2Sign: Post-cursor 2 Sign. Read-write. Cold-reset: 0. 1=Increases output voltage strength. 0=Lowers output voltage strength.
2:0	Reserved.

D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5 Phy Transmit Link Configuration

Table 105: Recommended link configuration

Link configuration	GangedModeEn	IsOwnMstr
x1 (1 lane per sublink)	0	1
x2 (2 lanes per sublink)	0	1
x4 (4 lanes per sublink)	0	0
x8	1	0

Table 106: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		6385h	6305h	6285h	6205h	6185h	6105h	6085h	6005h
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8
P_DP0_	0122h	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0	-	-	-	-

Table 106: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5

P_DP1_	0122h	-	-	-	-	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_DP2_	0123h	-	TX[P,N]6	TX[P,N]5	TX[P,N]4	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0

Table 107: **Broadcast Mapping** for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5

D0F0xE0[31:16]	D0F0xE0[15:0]		
	7705h	7605h	7005h
0120h	D0F0xE4_x0120_6[3:2][8,0]5	D0F0xE4_x0120_6[1:0][8,0]5	D0F0xE4_x0120_6[3:0][8,0]5
0121h	D0F0xE4_x0121_6[3:2][8,0]5	D0F0xE4_x0121_6[1:0][8,0]5	D0F0xE4_x0121_6[3:0][8,0]5
0221h	D0F0xE4_x0221_6[3:2][8,0]5	D0F0xE4_x0221_6[1:0][8,0]5	D0F0xE4_x0221_6[3:0][8,0]5
0122h	D0F0xE4_x0122_6[3:2][8,0]5	D0F0xE4_x0122_6[1:0][8,0]5	D0F0xE4_x0122_6[3:0][8,0]5
0123h	D0F0xE4_x0123_6[3:2][8,0]5	D0F0xE4_x0123_6[1:0][8,0]5	D0F0xE4_x0123_6[3:0][8,0]5

Bits	Description
31	GangedModeEn. Read-write. Reset: 1. BIOS: Table 105 . 1=Enables link ganged mode. 0=Disables link ganged mode.
30	Reserved.
29	IsOwnMstr. Read-write. Reset: 0. BIOS: Table 105 . 1=Enables the lane to self initialize its own read pointer.
28:0	Reserved.

D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]6 Phy Transmit Nominal Deemphasis Control

This register specifies the deemphasis, or preemphasis settings in the case of display port mode, and voltage margining settings for the transmit drivers.

Table 108: Recommended preemphasis settings

Conditions			D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]6		
Link Type	Preemphasis	Peak-to-peak Voltage	DeemphGen1Nom	TxMarginNom	
Display Port	0dB	1.2V	0	0	
		0.8V	0	42	
		0.6V	0	64	
		0.4V	0	85	
	3.5dB	0.8V	42	42	0
		0.6V	32	32	32
		0.4V	21	21	64
	6dB	0.6V	64	64	0
		0.4V	42	42	42
	9.5dB	0.4V	85	85	0
PCIe HDMI	-	1.2V	42	0	
DVI	-	1.2V	11	0	

Table 109: Index Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]6

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		6386h	6306h	6286h	6206h	6186h	6106h	6086h	6006h
P_GPP_	0120h	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_GFX_	0121h	TX[P,N]7	TX[P,N]6	TX[P,N]5	TX[P,N]4	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_GFX_	0221h	TX[P,N]15	TX[P,N]14	TX[P,N]13	TX[P,N]12	TX[P,N]11	TX[P,N]10	TX[P,N]9	TX[P,N]8
P_DP0_	0122h	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0	-	-	-	-
P_DP1_	0122h	-	-	-	-	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0
P_DP2_	0123h	-	TX[P,N]6	TX[P,N]5	TX[P,N]4	TX[P,N]3	TX[P,N]2	TX[P,N]1	TX[P,N]0

Table 110: Broadcast Mapping for D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]6

D0F0xE0[31:16]	D0F0xE0[15:0]		
	7706h	7606h	7006h
0120h	D0F0xE4_x0120_6[3:2][8,0]6	D0F0xE4_x0120_6[1:0][8,0]6	D0F0xE4_x0120_6[3:0][8,0]6
0121h	D0F0xE4_x0121_6[3:2][8,0]6	D0F0xE4_x0121_6[1:0][8,0]6	D0F0xE4_x0121_6[3:0][8,0]6
0221h	D0F0xE4_x0221_6[3:2][8,0]6	D0F0xE4_x0221_6[1:0][8,0]6	D0F0xE4_x0221_6[3:0][8,0]6
0122h	D0F0xE4_x0122_6[3:2][8,0]6	D0F0xE4_x0122_6[1:0][8,0]6	D0F0xE4_x0122_6[3:0][8,0]6
0123h	D0F0xE4_x0123_6[3:2][8,0]6	D0F0xE4_x0123_6[1:0][8,0]6	D0F0xE4_x0123_6[3:0][8,0]6

Bits	Description
31:16	Reserved.
15:8	DeemphGen1Nom. Read-write. Reset: 42. BIOS: Table 108 . This field specifies the post cursor deemphasis setting. Value must be less than or equal to 104.
7:0	TxMarginNom. Read-write. Reset: 0. BIOS: Table 108 . This field specifies the voltage margining setting of the transmit driver. Value must be less than or equal to 104.

D0F0xE4_x0[2:1]2[1:0]_[D:C][7:0][8,0]5 Termination Mode Control

Table 111: Index Mapping for D0F0xE4_x0[2:1]2[1:0]_[D:C][7:0][8,0]5

Pin Group	D0F0xE0[31:16]	D0F0xE0[15:0]							
		C385h	C305h	C285h	C205h	C185h	C105h	C085h	C005h
P_GPP_	0120h	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0	-	-	-	-
P_UMI_	0120h	-	-	-	-	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0121h	RX[P,N]7	RX[P,N]6	RX[P,N]5	RX[P,N]4	RX[P,N]3	RX[P,N]2	RX[P,N]1	RX[P,N]0
P_GFX_	0221h	RX[P,N]15	RX[P,N]14	RX[P,N]13	RX[P,N]12	RX[P,N]11	RX[P,N]10	RX[P,N]9	RX[P,N]8

Table 112: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_[D:C][7:0][8,0]5

D0F0xE0[31:16]	D0F0xE0[15:0]		
	D705h	D605h	D005h

Table 112: Broadcast Mapping for D0F0xE4_x0[2:1]2[1:0]_ID:C][7:0][8,0]5

0120h	D0F0xE4_x0120_C[3:2][8,0]5	D0F0xE4_x0120_C[1:0][8,0]5	D0F0xE4_x0120_C[3:0][8,0]5
0121h	D0F0xE4_x0121_C[3:2][8,0]5	D0F0xE4_x0121_C[1:0][8,0]5	D0F0xE4_x0121_C[3:0][8,0]5
0221h	D0F0xE4_x0221_C[3:2][8,0]5	D0F0xE4_x0221_C[1:0][8,0]5	D0F0xE4_x0221_C[3:0][8,0]5

Bits	Description										
31:2	Reserved.										
1:0	TermMode. Read-write; updated-by-hardware. Cold-reset: 10b. Receiver termination mode. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>100 Ohm differential DC, keepers on.</td> </tr> <tr> <td>01b</td> <td>50 Ohm single ended to ground.</td> </tr> <tr> <td>10b</td> <td>High impedance</td> </tr> <tr> <td>11b</td> <td>100 ohm differential DC, keepers off.</td> </tr> </tbody> </table>	Bits	Description	00b	100 Ohm differential DC, keepers on.	01b	50 Ohm single ended to ground.	10b	High impedance	11b	100 ohm differential DC, keepers off.
Bits	Description										
00b	100 Ohm differential DC, keepers on.										
01b	50 Ohm single ended to ground.										
10b	High impedance										
11b	100 ohm differential DC, keepers off.										

3.3.4 Wrapper Registers

Table 113: Mapping for wrapper registers

D0F0xE0[31:16]	Wrapper
0130h	GPPSB
0131h	Gfx
0132h	DDI
0133h	DDI2

D0F0xE4_x013[1:0]_0046 Subsystem and Vendor ID

Bits	Description
31:16	SubsystemID: subsystem id. Read-write. Reset: 1234h. Specifies the value returned by D[8:2]F0xB4[SubsystemID] .
15:0	SubsystemVendorID: subsystem vendor id . Read-write. Reset: 1022h. Specifies the value returned by D[8:2]F0xB4[SubsystemVendorID] .

D0F0xE4_x013[1:0]_0080 Link Configuration

Bits	Description																				
31:4	Reserved.																				
3:0	StrapBifLinkConfig. Read-write; strap. Reset: Product-specific. BIOS: See Table 49 and Table 53 . <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>x16 IO Link (Gfx Only)</td> <td>0100b</td> <td>4 x1 IO Links (GPPFCH Only)</td> </tr> <tr> <td>0001b</td> <td>x4 IO Link (GPPFCH Only)</td> <td>0101b</td> <td>2 x8 IO Links (Gfx Only)</td> </tr> <tr> <td>0010b</td> <td>2 x2 IO Links (GPPFCH Only)</td> <td>011xb</td> <td>Reserved</td> </tr> <tr> <td>0011b</td> <td>1 x2 IO Link, 2 x1 IO Links (GPPFCH Only)</td> <td>1xxx</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	0000b	x16 IO Link (Gfx Only)	0100b	4 x1 IO Links (GPPFCH Only)	0001b	x4 IO Link (GPPFCH Only)	0101b	2 x8 IO Links (Gfx Only)	0010b	2 x2 IO Links (GPPFCH Only)	011xb	Reserved	0011b	1 x2 IO Link, 2 x1 IO Links (GPPFCH Only)	1xxx	Reserved
Bits	Definition	Bits	Definition																		
0000b	x16 IO Link (Gfx Only)	0100b	4 x1 IO Links (GPPFCH Only)																		
0001b	x4 IO Link (GPPFCH Only)	0101b	2 x8 IO Links (Gfx Only)																		
0010b	2 x2 IO Links (GPPFCH Only)	011xb	Reserved																		
0011b	1 x2 IO Link, 2 x1 IO Links (GPPFCH Only)	1xxx	Reserved																		

D0F0xE4_x013[1:0]_0[C:8]00 Link Training Control**Table 114: Index address mapping for D0F0xE4_x013[1:0]_0[C:8]00**

Index	Function	Reset	Index	Function	Reset
0130_0800h	GPPFCH Port A	0000_0000h	0130_0C00h	GPPFCH Port E	0000_0001h
0130_0900h	GPPFCH Port B	0000_0001h	0131_0800h	Gfx Port A	0000_0001h
0130_0A00h	GPPFCH Port C	0000_0001h	0131_0900h	Gfx Port B	0000_0001h
0130_0B00h	GPPFCH Port D	0000_0001h	-	-	-

Bits	Description
31:1	Reserved.
0	HoldTraining: hold link training. Read-write. 1=Hold training on link.

D0F0xE4_x013[1:0]_0[C:8]03 Link Deemphasis Control**Table 115: Index address mapping for D0F0xE4_x013[1:0]_0[C:8]03**

Index	Function	Index	Function
0130_0803h	GPPFCH Port A	0130_0C03h	GPPFCH Port E
0130_0903h	GPPFCH Port B	0131_0803h	Gfx Port A
0130_0A03h	GPPFCH Port C	0131_0903h	Gfx Port B
0130_0B03h	GPPFCH Port D	-	-

Bits	Description
31:6	Reserved.
5	StrapBifDeemphasisSel. Read-write; strap. Reset: 1. Controls the default value of D[8:2]F0x88[SelectableDeemphasis] .
4:0	Reserved.

D0F0xE4_x013[3:0]_8011 Link Transmit Clock Gating Control

Bits	Description
31	StrapBifValid. Read-write. IF (REG==D0F0xE4_x0130_8011) THEN Reset: 0. ELSE Reset: 1. ENDIF. BIOS: 2.11.4.3.1. 0=Straps are latched. 1=Straps are not latched.
30:26	Reserved.
25	DdiDualLinkOverride. Read-write. Reset: 0. 1=Dual link DDI is clocked as a single link.
24	TxclkLcntGateEnable. Read-write. Reset: 0. BIOS: 1. 1=Enable clock gating the lane counter.
23	DebugBusClkEnable. Read-write. Reset: 0. 1=Enable the debug bus clock.
22:17	TxclkPermGateLatency. Read-write. Reset: 3Fh. Specifies the number of clocks to wait after detecting an entry into L1 before gating off the permanent clock branches.
16	RcvrDetClkEnable. Read-write. Reset: 0. 1=Enable the receiver detect clock.

15:10	TxclkRegsGateLatency. Read-write. Reset: 3Fh. Specifies the number of clocks to wait after idle is signalled before gating off the register clock branch.
9	TxclkRegsGateEnable. Read-write. Reset: 0. BIOS: 1. 1=Enable clock gating the register clock.
8	TxclkPermStop. Read-write. Reset: 0. 1=All transmitter clocks disabled. This bit should only be set if all links associated with the PCIe core are unconnected.
7	TxclkDynGateEnable. Read-write. Reset: 0. BIOS: 1. 1=Dynamic clock gating enabled. 0=Dynamic clock gating disabled.
6	TxclkPermGateEven. Read-write. Reset: 1. 1=Gate the permanent clock branches for an even number of clocks.
5:0	TxclkDynGateLatency. Read-write. Reset: 3Fh. Specifies the number of clocks to wait after idle is signalled before gating off the dynamic clock branch.

D0F0xE4_x013[3:0]_8012 Link Idle-Resume Clock Gating Control

Bits	Description
31:30	Reserved.
29:24	Pif2p5xIdleResumeLatency. Read-write. Reset: 0_0111b. Specifies the number of clocks to wait after enabling TXCLK2P5X_PIF before sending the acknowledge.
23	Pif2p5xIdleGateEnable. Read-write. Reset: 0. BIOS: 1. 1=Enable idle resume gating of TXCLK2P5X_PIF.
22	Reserved.
21:16	Pif2p5xIdleGateLatency. Read-write. Reset: 0_0001b. Specifies the number of clocks to wait before turning off TXCLK2P5X_PIF.
15:14	Reserved.
13:8	Pif1xIdleResumeLatency. Read-write. Reset: 0_0111b. Specifies the number of clocks to wait after enabling TXCLK1X_PIF before sending the acknowledge.
7	Pif1xIdleGateEnable. Read-write. Reset: 0. BIOS: 1. 1=Enable idle resume gating of TXCLK1X_PIF.
6	Reserved.
5:0	Pif1xIdleGateLatency. Read-write. Reset: 0_0001b. Specifies the number of clocks to wait before turning off TXCLK1X_PIF.

D0F0xE4_x013[3:0]_8013 Transmit Clock Pll Control

Reset: 0000_0001h.

Table 116: Reserved field mappings for D0F0xE4_x013[3:0]_8013

Register	Bits						
	20:13	12:11	10	7:6	5	3:2	1
D0F0xE4_x0130_8013	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved
D0F0xE4_x0131_8013	-	-	-	-	-	-	-
D0F0xE4_x0132_8013	-	Reserved	-	Reserved	-	Reserved	-
D0F0xE4_x0133_8013	-	Reserved	-	Reserved	-	Reserved	-

Bits	Description												
31:21	Reserved.												
20	TxclkSelDigBOverride . Read-write. 1=Override TxclkDigB selection.												
19:17	<p>TxclkSelDigB. Read-write. Specifies the source of the dig B clock when TxclkSelDigBOverride=1.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0x0b</td> <td>Phy Clock A</td> <td>10xb</td> <td>Phy Clock C</td> </tr> <tr> <td>0x1b</td> <td>Phy Clock B</td> <td>11xb</td> <td>Phy Clock D</td> </tr> </tbody> </table> <p>The selected clock with not function correctly if the divider logic is not enabled for the clock. The divider is enabled if (ClkDividerResetOverrideX MasterPciePllX D0F0xE4_x013[3:1]_804[3:0][OwnSlice]) =1.</p>	Bits	Definition	Bits	Definition	0x0b	Phy Clock A	10xb	Phy Clock C	0x1b	Phy Clock B	11xb	Phy Clock D
Bits	Definition	Bits	Definition										
0x0b	Phy Clock A	10xb	Phy Clock C										
0x1b	Phy Clock B	11xb	Phy Clock D										
16	TxclkSelDigAOverride . Read-write. 1=Override TxclkDigA selection.												
15:13	<p>TxclkSelDigA. Read-write. Specifies the source of the dig A clock when TxclkSelDigAOverride=1.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0x0b</td> <td>Phy Clock A</td> <td>10xb</td> <td>Phy Clock C</td> </tr> <tr> <td>0x1b</td> <td>Phy Clock B</td> <td>11xb</td> <td>Phy Clock D</td> </tr> </tbody> </table> <p>The selected clock with not function correctly if the divider logic is not enabled for the clock. The divider is enabled if (ClkDividerResetOverrideX MasterPciePllX D0F0xE4_x013[3:1]_804[3:0][OwnSlice]) =1.</p>	Bits	Definition	Bits	Definition	0x0b	Phy Clock A	10xb	Phy Clock C	0x1b	Phy Clock B	11xb	Phy Clock D
Bits	Definition	Bits	Definition										
0x0b	Phy Clock A	10xb	Phy Clock C										
0x1b	Phy Clock B	11xb	Phy Clock D										
12	TxclkSelPifDOverride . Read-write. 1=Override TxclkPifD selection.												
11	TxclkSelPifCOverride . Read-write. 1=Override TxclkPifC selection.												
10	TxclkSelPifBOverride . Read-write. 1=Override TxclkPifB selection.												
9	TxclkSelPifAOverride . Read-write. 1=Override TxclkPifA selection.												
8	TxclkSelCoreOverride . Read-write. 1=Override TxclkCore selection.												
7	ClkDividerResetOverrideD . Read-write. 1=Force clock divider D enabled.												
6	ClkDividerResetOverrideC . Read-write. 1=Force clock divider C enabled.												
5	ClkDividerResetOverrideB . Read-write. 1=Force clock divider B enabled.												
4	ClkDividerResetOverrideA . Read-write. 1=Force clock divider A enabled.												
3	MasterPciePllD . Read-write. 1=Pll D is the master source for all PCIe transmitter clock branches.												
2	MasterPciePllC . Read-write. 1=Pll C is the master source for all PCIe transmitter clock branches.												
1	MasterPciePllB . Read-write. 1=Pll B is the master source for all PCIe transmitter clock branches.												
0	MasterPciePllA . Read-write. 1=Pll A is the master source for all PCIe transmitter clock branches.												

D0F0xE4_x013[3:0]_8014 Link Transmit Clock Gating Control 2

Reset: 0000_0000h.

Table 117: Reserved field mappings for D0F0xE4_x013[3:0]_8014

Register	Bits				
	27:26	19:18	15:14	11:10	5:4
D0F0xE4_x0130_8014	Reserved	Reserved	Reserved	Reserved	Reserved
D0F0xE4_x0131_8014	-	-	-	-	-
D0F0xE4_x0132_8014	Reserved	Reserved	Reserved	Reserved	Reserved
D0F0xE4_x0133_8014	Reserved	Reserved	Reserved	Reserved	Reserved

Bits	Description
31:28	Reserved.
27	DdiGateDigDEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the dig d clock branches in DDI mode.
26	DdiGateDigCEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the dig c clock branches in DDI mode.
25	DdiGateDigBEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the dig b clock branches in DDI mode.
24	DdiGateDigAEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the dig a clock branches in DDI mode.
23:21	Reserved.
20	TxclkPermGateOnlyWhenPlIPwrDn. Read-write. BIOS: 1. 1=Gating of the permanent clock branch only occurs when the PLL is powered down.
19	PcieGatePifD2p5xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF D 2.5x clock branches in PCIe mode.
18	PcieGatePifC2p5xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF C 2.5x clock branches in PCIe mode.
17	PcieGatePifB2p5xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF B 2.5x clock branches in PCIe mode.
16	PcieGatePifA2p5xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF A 2.5x clock branches in PCIe mode.
15	PcieGatePifD1xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF D 1x clock branches in PCIe mode.
14	PcieGatePifC1xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF C 1x clock branches in PCIe mode.
13	PcieGatePifB1xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF B 1x clock branches in PCIe mode.
12	PcieGatePifA1xEnable. Read-write. BIOS: 1. 1=Enable gating of the PIF A 1x clock branches in PCIe mode.
11	DdiGatePifD2p5xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF D 2.5x clock branches in DDI mode.
10	DdiGatePifC2p5xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF C 2.5x clock branches in DDI mode.
9	DdiGatePifB2p5xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF B 2.5x clock branches in DDI mode.
8	DdiGatePifA2p5xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF A 2.5x clock branches in DDI mode.
7:6	Reserved.
5	DdiGatePifD1xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF D 1x clock branches in DDI mode.
4	DdiGatePifC1xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF C 1x clock branches in DDI mode.

3	DdiGatePifB1xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF B 1x clock branches in DDI mode.
2	DdiGatePifA1xEnable. Read-write. BIOS: IF (REG==D0F0xE4_x0130_8014) THEN 0 ELSE 1 ENDIF. 1=Enable gating of the PIF A 1x clock branches in DDI mode.
1	TxclkPrbsGateEnable. Read-write. BIOS: 1. 1=Enable gating of the PRBS clock branch.
0	TxclkPermGateEnable. Read-write. BIOS: 1. 1=Enable gating of the permanent clock branch.

D0F0xE4_x013[3:0]_8015 IO Link IOC Control

Bits	Description
31	RefclkBphyGateEnable. Read-write. Reset:0. BIOS: 1. 1=Enable gating of REFCLK_BPHY.
30	Reserved.
29:24	RefclkBphyGateLatency. Read-write. Reset:3Fh. BIOS: 0. Specifies the number of clocks to wait before turning of f REFCLK_BPHY.
23	RefclkRegsGateEnable. Read-write. Reset:0. BIOS: 1. 1=Enable gating of REFCLK_REGS.
22	Reserved.
21:16	RefclkRegsGateLatency. Read-write. Reset:3Fh. Specifies the number of clocks to wait before turning of f REFCLK_REGS.
15:0	Reserved.

D0F0xE4_x013[3:0]_8016 Link Clock Switching Control

Reset: 003F_001Fh.

Table 118: Reserved field mappings for D0F0xE4_x013[3:0]_8016

Register	Bits	
	11:10	9
D0F0xE4_x0130_8016	Reserved	Reserved
D0F0xE4_x0131_8016	-	-
D0F0xE4_x0132_8016	Reserved	-
D0F0xE4_x0133_8016	Reserved	-

Bits	Description
31:24	Reserved.
23	LclkDynGateEnable. Read-write. BIOS: 1. 1=Enable LCLK_DYN clock gating.
22	LclkGateFree. Read-write. BIOS: 1. 1=LCLK gating is controlled independent of TXCLK gating.
21:16	LclkDynGateLatency. Read-write. Specifies the number of clocks to wait before turning off LCLK_DYN.
15:6	Reserved.
5:0	CalibAckLatency. Read-write. BIOS: IF (REG== D0F0xE4_x0130_8016) THEN 0 ELSE 1Fh. ENDIF. Specifies the number of clocks after calibration is complete before the acknowledge signal is asserted.

D0F0xE4_x013[3:0]_8020 Lane Control

Reset: 0000_0000h.

Bits	Description
31:4	Reserved.
3	PrbsPcieLbSelect . Read-write. 1=Loopback data selected. 0=IO/PRBS data selected.
2:0	Reserved.

D0F0xE4_x013[3:0]_8021 Transmitter Lane Mux

Reset: 7654_3210h.

Bits	Description																								
31:28	Lanes1514 . Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 15 and 14. See: Lanes10.																								
27:24	Lanes1312 . Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 13 and 12. See: Lanes10.																								
23:20	Lanes1110 . Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 11 and 10. See: Lanes10.																								
19:16	Lanes98 . Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 9 and 8. See: Lanes10.																								
15:12	Lanes76 . Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 7 and 6. See: Lanes10.																								
11:8	Lanes54 . Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 5 and 4. See: Lanes10.																								
7:4	Lanes32 . Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 3 and 2. See: Lanes10.																								
3:0	<p>Lanes10. Read-write. Specifies the controller lanes that are mapped to PIF TX lanes 1 and 0.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Controller lanes 1 and 0</td> <td>5h</td> <td>Controller lanes 11 and 10.</td> </tr> <tr> <td>1h</td> <td>Controller lanes 3 and 2.</td> <td>6h</td> <td>Controller lanes 13 and 12.</td> </tr> <tr> <td>2h</td> <td>Controller lanes 5 and 4.</td> <td>7h</td> <td>Controller lanes 15 and 14.</td> </tr> <tr> <td>3h</td> <td>Controller lanes 7 and 6.</td> <td>Fh-8h</td> <td>Reserved.</td> </tr> <tr> <td>4h</td> <td>Controller lanes 9 and 8.</td> <td></td> <td></td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	0h	Controller lanes 1 and 0	5h	Controller lanes 11 and 10.	1h	Controller lanes 3 and 2.	6h	Controller lanes 13 and 12.	2h	Controller lanes 5 and 4.	7h	Controller lanes 15 and 14.	3h	Controller lanes 7 and 6.	Fh-8h	Reserved.	4h	Controller lanes 9 and 8.		
Bits	Definition	Bits	Definition																						
0h	Controller lanes 1 and 0	5h	Controller lanes 11 and 10.																						
1h	Controller lanes 3 and 2.	6h	Controller lanes 13 and 12.																						
2h	Controller lanes 5 and 4.	7h	Controller lanes 15 and 14.																						
3h	Controller lanes 7 and 6.	Fh-8h	Reserved.																						
4h	Controller lanes 9 and 8.																								

D0F0xE4_x013[3:0]_8022 Receiver Lane Mux

Reset: 7654_3210h.

Bits	Description
31:28	Lanes1514 . Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 15 and 14. See: Lanes10.
27:24	Lanes1312 . Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 13 and 12. See: Lanes10.

23:20	Lanes1110. Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 11 and 10. See: Lanes10.																								
19:16	Lanes98. Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 9 and 8. See: Lanes10.																								
15:12	Lanes76. Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 7 and 6. See: Lanes10.																								
11:8	Lanes54. Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 5 and 4. See: Lanes10.																								
7:4	Lanes32. Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 3 and 2. See: Lanes10.																								
3:0	<p>Lanes10. Read-write. Specifies the PIF RX lanes that are mapped to controller lanes 1 and 0.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>PIF RX lanes 1 and 0</td> <td>5h</td> <td>PIF RX lanes 11 and 10.</td> </tr> <tr> <td>1h</td> <td>PIF RX lanes 3 and 2.</td> <td>6h</td> <td>PIF RX lanes 13 and 12.</td> </tr> <tr> <td>2h</td> <td>PIF RX lanes 5 and 4.</td> <td>7h</td> <td>PIF RX lanes 15 and 14.</td> </tr> <tr> <td>3h</td> <td>PIF RX lanes 7 and 6.</td> <td>Fh-8h</td> <td>Reserved.</td> </tr> <tr> <td>4h</td> <td>PIF RX lanes 9 and 8.</td> <td></td> <td></td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	0h	PIF RX lanes 1 and 0	5h	PIF RX lanes 11 and 10.	1h	PIF RX lanes 3 and 2.	6h	PIF RX lanes 13 and 12.	2h	PIF RX lanes 5 and 4.	7h	PIF RX lanes 15 and 14.	3h	PIF RX lanes 7 and 6.	Fh-8h	Reserved.	4h	PIF RX lanes 9 and 8.		
Bits	Definition	Bits	Definition																						
0h	PIF RX lanes 1 and 0	5h	PIF RX lanes 11 and 10.																						
1h	PIF RX lanes 3 and 2.	6h	PIF RX lanes 13 and 12.																						
2h	PIF RX lanes 5 and 4.	7h	PIF RX lanes 15 and 14.																						
3h	PIF RX lanes 7 and 6.	Fh-8h	Reserved.																						
4h	PIF RX lanes 9 and 8.																								

D0F0xE4_x013[3:0]_8023 Lane Enable

Reset: 0000_FFFFh.

Bits	Description				
31:16	Reserved.				
15:0	<p>LaneEnable. Read-write. 1=Lane enabled for transmit.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>[15:0]</td> <td>Lane <BIT> enable</td> </tr> </tbody> </table>	Bit	Definition	[15:0]	Lane <BIT> enable
Bit	Definition				
[15:0]	Lane <BIT> enable				

D0F0xE4_x013[3:0]_8025 Lane Mux Power Sequence Control

Reset: 1F1F_1F1Fh.

Bits	Description
31:30	Reserved.
29	LMLinkSpeed3. Read-write. 1=5GHz. 0=2.5GHz.
28:27	LMRxPhyCmd3. Read-write. Specifies the receiver state for lanes 15-12. See: LMRxPhyCmd0.
26:24	LMTxPhyCmd3. Read-write. Specifies the transmitter state for lanes 15-12. See: LMTxPhyCmd0.
23:22	Reserved.
21	LMLinkSpeed2. Read-write. 1=5GHz. 0=2.5GHz.
20:19	LMRxPhyCmd2. Read-write. Specifies the receiver state for lanes 11-8. See: LMRxPhyCmd0.
18:16	LMTxPhyCmd2. Read-write. Specifies the transmitter state for lanes 11-8. See: LMTxPhyCmd0.
15:14	Reserved.
13	LMLinkSpeed1. Read-write. 1=5GHz. 0=2.5GHz.
12:11	LMRxPhyCmd1. Read-write. Specifies the receiver state for lanes 7-4. See: LMRxPhyCmd0.
10:8	LMTxPhyCmd1. Read-write. Specifies the transmitter state for lanes 7-4. See: LMTxPhyCmd0.

7:6	Reserved.			
5	LMLinkSpeed0. Read-write. 1=5GHz. 0=2.5GHz.			
4:3	LMRxPhyCmd0. Read-write. Specifies the receiver state for lanes 3-0.			
	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>
	00b	On.	10b	Standby 2 (L1).
	01b	Standby 1 (L0s).	11b	Off.
2:0	LMTxPhyCmd0. Read-write. Specifies the transmitter state for lanes 3-0.			
	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>
	000b	On.	100b	Receiver detect.
	001b	Standby 1 (L0s).	101b	Reserved.
	010b	Standby 2 (L1).	110b	Reserved.
	011b	Reserved.	111b	Off.

D0F0xE4_x013[1:0]_8031 Lane Counter Status

Reset: 0000_0000h.

Bits	Description
31:17	Reserved.
16	LnCntValid. Read-write. 1=LnCntBandwidth contains a valid bandwidth measurement.
15:10	Reserved.
9:0	LnCntBandwidth. Read-write. Estimated lane bandwidth in 10 MB/s units.

D0F0xE4_x013[3:0]_8060 Soft Reset Command 0

Cold reset: 0000_0000h.

Bits	Description
31:18	Reserved.
17	Bif0CalibrationReset. Read-write. 1=The BIF 0 calibration block reset is asserted.
16	Bif0GlobalReset. Read-write. 1=The BIF 0 global reset is asserted.
15:3	Reserved.
2	ResetComplete. Read-only. 1=Reset cycle is complete.
1	Reserved.
0	Reconfigure. Read-write; Cleared-when-done. 1=Trigger atomic reconfiguration if D0F0xE4_x013[3:0]_8062[ReconfigureEn]=1 .

D0F0xE4_x013[3:0]_8062 Soft Reset Control 0

Cold reset: 0001_0800h.

Bits	Description
31:12	Reserved.
11	ConfigXferMode. Read-write. 1=PCIe core strap settings take effect immediately. 0=PCIe core strap settings take effect when the PCIe core is reset.

10	BlockOnIdle. Read-write. 1=The PCIe core must be idle before hardware initiates a reconfiguration. 0=The PCIe core does not have to be idle before hardware initiates a reconfiguration.
9:5	Reserved.
4:2	ResetPeriod. Read-write. BIOS:0. Specifies the amount of time that resets are asserted during a reconfiguration.
1	Reserved.
0	ReconfigureEn. Read-write. 1=Atomic reconfiguration enabled.

D0F0xE4_x013[3:0]_8063 Soft Reset Control 1

Cold reset: 0000_0C00h.

Bits	Description
31:15	Reserved.
14	ResetSrbmDcEn. Read-write. 1=The reset to the dc register interface block is asserted during an atomic reset or reconfiguration.
13	ResetSrbmGfxEn. Read-write. 1=The reset to the gfx register interface block is asserted during an atomic reset or reconfiguration.
12	ResetSrbmNbEn. Read-write. 1=The reset to the nb register interface block is asserted during an atomic reset or reconfiguration.
11:6	Reserved.
5	ResetSrbmIEn. Read-write. 1=The reset to the PCIe® core register interface block is applied during an atomic reset or reconfiguration.
4	ResetSrbm0En. Read-write. 1=The reset to the wrapper register interface block is applied during an atomic reset or reconfiguration.
3:0	Reserved.

D0F0xE4_x0130_80F0 BIOS Timer

Reset: 0000_0000h.

Bits	Description
31:0	MicroSeconds. Read-write; updated-by-hardware. This field increments once every microsecond when the timer is enabled. The counter will roll over and continue counting when it reaches its FFFF_FFFFh. A write to this register causes the counter to reset and begin counting from the value written.

D0F0xE4_x0130_80F1 BIOS Timer Control

Reset: 0000_0064h.

Bits	Description
31:8	Reserved.
7:0	ClockRate. Read-write. Specifies the frequency of the reference clock in 1 MHz increments. <u>Bits</u> <u>Definition</u> 00h Timer disabled FFh-01h <ClockRate> MHz

D0F0xE4_x013[3:1]_804[3:0] DDI Slice

Reset: 0000_0000h.

Table 119: Index address mapping for D0F0xE4_x013[3:1]_804[3:0]

D0F0xE0[31:16]	D0F0xE0[15:0]			
	8040h	8041h	8042h	8043h
0131h	Gfx Lanes 0-3	Gfx Lanes 4-7	Gfx Lanes 8-11	Gfx Lanes 12-15
0132h	DDI Lanes 0-3	DDI Lanes 4-7	DDI Lanes 8-11	DDI Lanes 12-15
0133h	DDI2 Lanes 0-3	DDI2 Lanes 4-7	DDI2 Lanes 8-11	DDI2 Lanes 12-15

Bits	Description
31:1	Reserved.
0	OwnSlice. Read-write. 1=DDI asserts control over the PCIe lanes specified in Table 119 .

D0F0xE4_x013[3:1]_804[E:8] DDI Dig

Reset: 0000_0700h.

Table 120: Register Stream mappings for D0F0xE4_x013[3:1]_804[E:8]

D0F0xE4_x013[3:1]_804[E:8]	Function
D0F0xE4_x013[3:1]_8048	Stream A
D0F0xE4_x013[3:1]_8049	Stream B
D0F0xE4_x013[3:1]_804A	Stream C
D0F0xE4_x013[3:1]_804B	Stream D
D0F0xE4_x013[3:1]_804C	Stream E
D0F0xE4_x013[3:1]_804D	Stream F
D0F0xE4_x013[3:1]_804E	Stream G

Bits	Description
31:26	Reserved.
25	CntDig. Read-write. 1=Software asserts control over Dig TxPhyCmd and LinkSpeed.
24	CntPhy. Read-write. 1=Software asserts control over the phy state machine.

23	Reserved.																				
22	Nxt_Lnsprd. Read-write. 1=Set the value for the dig link speed in case of an override.																				
21:19	Nxt_phycmd. Read-write. 1=Set the value for the dig Tx phy command in case of an override.																				
18:16	Nxt_State. Read-write. This specifies the next state for the DDI FSM.																				
15:11	Reserved.																				
10:8	<p>Dig_pwrdn_value. Read-write. Specifies the phy power state for links associated with dig streams that are not enabled.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>On.</td> <td>100b</td> <td>Receiver detect.</td> </tr> <tr> <td>001b</td> <td>Standby 1 (L0s).</td> <td>101b</td> <td>Reserved.</td> </tr> <tr> <td>010b</td> <td>Standby 2 (L1).</td> <td>110b</td> <td>Reserved.</td> </tr> <tr> <td>011b</td> <td>Reserved.</td> <td>111b</td> <td>Off.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	000b	On.	100b	Receiver detect.	001b	Standby 1 (L0s).	101b	Reserved.	010b	Standby 2 (L1).	110b	Reserved.	011b	Reserved.	111b	Off.
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>																		
000b	On.	100b	Receiver detect.																		
001b	Standby 1 (L0s).	101b	Reserved.																		
010b	Standby 2 (L1).	110b	Reserved.																		
011b	Reserved.	111b	Off.																		
7	Reserved.																				
6	Hbr2Support. Read-write. 1=HBR is supported.																				
5	Reserved.																				
4	Hbr2Active. Read-write. 1=If (Hbr2Support==1) THEN HBR2 will be enabled.																				
3	Reserved.																				
2	PwrDnCpl. Read-only; updated-by-hardware. 1=PHY state machine is in the powered up state.																				
1	Reserved.																				
0	PwrDnCpl. Read-only; updated-by-hardware. 1=PHY state machine is in the power off state.																				

3.4 Device 0 Function 2 (IOMMU) Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space]. See 2.12.1 [IOMMU Configuration Space].

D0F2x00 Device/Vendor ID

Bits	Description
31:16	DeviceId. Read-only. Reset: 1419h.
15:0	VendorId. Read-only. Reset: 1022h.

D0F2x04 Status/Command

Bits	Description
31	ParityErrorDetected. Read; write-1-to-clear. Reset: 0.
30	SignaledSystemError. Read-only. Reset: 0.
29	ReceivedMasterAbort. Read; write-1-to-clear. Reset: 0.
28	ReceivedTargetAbort. Read; write-1-to-clear. Reset: 0.
27	SignalTargetAbort. Read-only. Reset: 0.
26:25	Reserved.
24	MasterDataError. Read; write-1-to-clear. Reset: 0.
23:21	Reserved.
20	CapList. Read-only. Reset: 1. 1=Capability list supported.
19	IntStatus. Read-only. Reset: 0. 1=INTx message pending.
18:11	Reserved.
10	InterruptDis. Read-write. Reset: 0. 1=INTx interrupt message generation disabled.
9	Reserved.
8	SerrEn. Read-only. Reset: 0. 1=Enables reporting of non-fatal and fatal errors detected.
7	Reserved.
6	ParityErrorEn. Read-write. Reset: 0. 1=Enables setting of ParityErrorDetected status bit.
5:3	Reserved.
2	BusMasterEn. Read-write. Reset: 0. 1=Enables DMA request generation.
1	MemAccessEn. Read-only. Reset: 0.
0	IoAccessEn. Read-only. Reset: 0.

D0F2x08 Class Code/Revision ID

Reset: 0806_00xxh.

Bits	Description
31:8	ClassCode: class code. Read-only. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only.

D0F2x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:24	BIST. Read-only.
23:16	HeaderTypeReg. Read-only. 80h=Type 0 multi-function device.
15:8	LatencyTimer. Read-write.
7:0	CacheLineSize. Read-only.

D0F2x2C Subsystem and Subvendor ID

Bits	Description
31:16	SubsystemId. Read-only. Reset: 0.
15:0	SubsystemVendorId. Read-only. Reset: 0.

D0F2x34 Capabilities Pointer

Bits	Description
31:8	Reserved.
7:0	CapPtr. Read-only. Reset: 40h.

D0F2x3C Interrupt Line

Bits	Description
31:16	Reserved.

15:8	InterruptPin. Read-only. Reset: 01h. This field indicates the INTx line used to generate legacy interrupts. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Reserved.</td> </tr> <tr> <td>01h</td> <td>INTA.</td> </tr> <tr> <td>02h</td> <td>INTB.</td> </tr> <tr> <td>03h</td> <td>INTC.</td> </tr> <tr> <td>04h</td> <td>INTD.</td> </tr> <tr> <td>FFh-05h</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	00h	Reserved.	01h	INTA.	02h	INTB.	03h	INTC.	04h	INTD.	FFh-05h	Reserved.
Bits	Description														
00h	Reserved.														
01h	INTA.														
02h	INTB.														
03h	INTC.														
04h	INTD.														
FFh-05h	Reserved.														
7:0	InterruptLine. Read-write. Reset: 0. This field is read/write for software compatibility. It controls no hardware.														

D0F2x40 IOMMU Capability

Bits	Description
31:28	Reserved.
27	IommuEfrSup. Read-only. Reset: 1. 1=Indicates IOMMUx30 [Extended Feature Low] is supported. 0= IOMMUx30 is reserved.
26	IommuNpCache. Read-only. Reset: 0. 1=Indicates that the IOMMU caches page table entries that are marked as not present. When this bit is set, software must issue an invalidate after any change to a PDE or PTE. 0=Indicates that the IOMMU caches only page table entries that are marked as present. When this bit is clear, software must issue an invalidate after any change to a PDE or PTE marked present before the change.
25	IommuHtTunnelSup. Read-only. Reset: 0.
24	IommuIoTlbsup. Read-only. Reset: 1. Indicates support for remote IOTLBs.
23:19	IommuCapRev. Read-only. Reset: 1. Specifies the IOMMU interface revision.
18:16	IommuCapType. Read-only. Reset: 3h. Specifies the layout of the Capability Block as an IOMMU capability block.
15:8	IommuCapPtr. Read-only. Reset: 54h. Indicates the location of the next capability block.
7:0	IommuCapId. Read-only. Reset: Fh. Indicates a Secure Device capability block.

D0F2x44 IOMMU Base Address Low

Bits	Description
31:14	IommuBaseAddr[31:14]: iommu base address bits[31:14]. IF (D0F2x44[IommuEnable]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0. IommuBaseAddr[63:14] = {D0F2x48[IommuBaseAddr[63:32]], IommuBaseAddr[31:14]}. IommuBaseAddr[63:14] specifies the base address of the IOMMU memory mapped control registers. In order to use the IOMMU event counters, IommuBaseAddr[18:14] must be 0_0000b.
13:1	Reserved.
0	IommuEnable. Read; write-1-only. Reset: 0. 1=IOMMU accepts memory accesses to the address specified in IommuBaseAddr[63:14]. When this bit is set, all IOMMU RW capability registers in PCI configuration space are locked.

D0F2x48 IOMMU Base Address High

Bits	Description
31:0	IommuBaseAddr[63:32]: iommu base address bits[63:32]. See: D0F2x44[IommuBase-Addr[31:14]] .

D0F2x4C IOMMU Range

Bits	Description
31:24	IommuLastDevice. Read-only. Reset: 0. Indicates device and function number of the last integrated device associated with the IOMMU.
23:16	IommuFirstDevice. Read-only. Reset: 0. Indicates device and function number of the first integrated device associated with the IOMMU.
15:8	IommuBusNumber. Read-only. Reset: 0. Indicates the bus number that IommuLastDevice and IommuFirstDevice reside on.
7	IommuRngValid. Read-only. Reset: 0. 1=The IommuBusNumber, IommuFirstDevice, and IommuLastDevice fields are valid. Although the register contents are valid, software is encouraged to use I/O topology information.. 0=Software must use I/O topology information.
6:5	Reserved.
4:0	IommuUnitId. Read-only. Reset: 0.

D0F2x50 IOMMU Miscellaneous Information Register

Bits	Description
31:27	IommuMsiNumPpr. Read-only. Reset: 0. This field must indicate which MSI vector is used for the interrupt message generated by the IOMMU for the peripheral page service request log when IOMMUx30[PprSup]=1 . This field must be 0 when IOMMUx30[PprSup]=0 . For MSI there can be only one IOMMU so this field must be 0. This interrupt is not remapped by the IOMMU.
26:23	Reserved.
22	IommuHtAtsResv. IF (D0F2x44[IommuEnable]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0. 1=The HyperTransport Address Translation address range for ATS responses is reserved and cannot be translated by the IOMMU. 0=The Address Translation address range can be translated by the IOMMU.

21:15	IommuVaSize. Read-only. Reset: 40h. This field must indicate the size of the maximum virtual address processed by the IOMMU. The value is the (unsigned) binary log of the maximum address size. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>19h-00h</td> <td>Reserved.</td> </tr> <tr> <td>20h</td> <td>32 bits.</td> </tr> <tr> <td>27h-21h</td> <td>Reserved.</td> </tr> <tr> <td>28h</td> <td>40 bits.</td> </tr> <tr> <td>2Fh-29h</td> <td>Reserved.</td> </tr> <tr> <td>30h</td> <td>48 bits.</td> </tr> <tr> <td>3Fh-31h</td> <td>Reserved.</td> </tr> <tr> <td>40h</td> <td>64 bits.</td> </tr> <tr> <td>7Fh-41h</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	19h-00h	Reserved.	20h	32 bits.	27h-21h	Reserved.	28h	40 bits.	2Fh-29h	Reserved.	30h	48 bits.	3Fh-31h	Reserved.	40h	64 bits.	7Fh-41h	Reserved.
<u>Bits</u>	<u>Description</u>																				
19h-00h	Reserved.																				
20h	32 bits.																				
27h-21h	Reserved.																				
28h	40 bits.																				
2Fh-29h	Reserved.																				
30h	48 bits.																				
3Fh-31h	Reserved.																				
40h	64 bits.																				
7Fh-41h	Reserved.																				
14:8	IommuPaSize. Read-only. Reset: 28h. This field must indicate the size of the maximum physical address generated by the IOMMU. The value is the (unsigned) binary log of the maximum address size. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>27h-00h</td> <td>Reserved.</td> </tr> <tr> <td>28h</td> <td>40 bits.</td> </tr> <tr> <td>7Fh-29h</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	27h-00h	Reserved.	28h	40 bits.	7Fh-29h	Reserved.												
<u>Bits</u>	<u>Description</u>																				
27h-00h	Reserved.																				
28h	40 bits.																				
7Fh-29h	Reserved.																				
7:5	IommuGvaSize. Read-only. Reset: 010b. Indicates the size of the maximum guest virtual address processed by the IOMMU. 010b=48 bits. All other values are reserved.																				
4:0	IommuMsiNum. Read-only. Reset: 0. Indicates the MSI vector used for interrupt messages generated by the IOMMU.																				

D0F2x54 IOMMU MSI Capability Register

Bits	Description
31:24	Reserved.
23	Msi64En. Read-only. Reset: 1. 1=64-bit MSI addressing is supported
22:20	MsiMultMessEn. Read-only. Reset: 0. Specifies the number of MSI messages assigned to this function.
19:17	MsiMultMessCap. Read-only. Reset: 0. Specifies the number of MSI messages requested by this function.
16	MsiEn. Read-write. Reset: 0. 1=Enables MSI for this function and causes legacy interrupts to be disabled.
15:8	MsiCapPtr. Read-only. Reset: 64h. Pointer to the next capability register offset.
7:0	MsiCapId. Read-only. Reset: 5h. Indicates that this is the MSI capability.

D0F2x58 IOMMU MSI Address Low

Bits	Description
31:2	MsiAddr[31:2]. Read-write. Reset: 0. This register specifies the lower address bits used to issue MSI messages.
1:0	Reserved.

D0F2x5C IOMMU MSI Address High

Bits	Description
31:0	MsiAddr[63:32] . Read-write. Reset: 0. This register specifies the upper address bits used to issue MSI messages.

D0F2x60 IOMMU MSI Data

Bits	Description
31:16	Reserved.
15:0	MsiData . Read-write. Reset: 0. This register specifies the data issued with MSI messages.

D0F2x64 IOMMU MSI Mapping Capability

Bits	Description
31:27	MsiMapCapType . Read-only. Reset: 15h. Indicates the MSI Mapping Capability.
26:18	Reserved.
17	MsiMapFixd . Read-only. Reset: 1. Always set to 1 to indicate that this device only maps MSI interrupts with address 0xFEEEx_xxxx onto Hypertransport interrupts and that the mapping range is not programmable
16	MsiMapEn . Read-only. Reset: 1. Always set to 1 to indicate that the MSI Mapping Capability is always enabled
15:8	MsiMapCapPtr . Read-only. Reset: 0. Points to the next capability list item
7:0	MsiMapCapId . Read-only. Reset: 8h. Indicates a Hypertransport capability list item

D0F2x6C IOMMU Control

Bits	Description
31:10	Reserved.
9	EfrSupW . Read-write. Reset: 1. This field sets the value of D0F2x40[EfrSup] .
8	IoTlbsupW . Read-write. Reset: 1. This field sets the value of D0F2x40[IommuIoTlbsup] .
7:4	MinorRevIdW . Read-write. Reset: 0h. This field sets the value of D0F2x08[RevID[3:0]] .
3	Reserved.
2:0	InterruptPinW . Read-write. Reset: 001b. This field sets the value of D0F2x3C[InterruptPin] .

D0F2x70 IOMMU MMIO Control Low

Bits	Description
31:12	Reserved.

11:10	HatsW . Read-write. Reset: 2h. This field sets the value of IOMMUx30[HATS] .
9	PcSupW . Read-write. Reset: 1. BIOS: 0. This field sets the value of IOMMUx30[PcSup] . This bit should not be set.
8	Reserved.
7	Reserved.
6	IaSupW . Read-write. Reset: 1. This field sets the value of IOMMUx30[IaSup] .
5	Reserved.
4	GtSupW . Read-write. Reset: 1. This field sets the value of IOMMUx30[GtSup] .
3	NxSupW . Read-write. Reset: 0. This field sets the value of IOMMUx30[NxSup] .
2	Reserved.
1	PprSupW . Read-write. Reset: 1. This field sets the value of IOMMUx30[PprSup] .
0	PrefSupW . Read-write. Reset: 1. This field sets the value of IOMMUx30[PrefSup] .

D0F2x74 IOMMU MMIO Control High

Bits	Description
31:4	Reserved.
3:0	PasMaxW . Read-write. Reset: 8h. This field sets the value of IOMMUx34[PasMax] .

D0F2x78 IOMMU Range Control

The fields in this register set the values of the corresponding fields in [D0F2x4C](#).

Bits	Description
31:24	LastDeviceW . Read-write. Reset: 0.
23:16	FirstDeviceW . Read-write. Reset: 0.
15:8	BusNumberW . Read-write. Reset: 0.
7	RngValidW . Read-write. Reset: 0.
6:0	Reserved.

D0F2xF0 IOMMU L2 Config Index

The index/data pair registers [D0F2xF0](#) and [D0F2xF4](#) is used to access the registers [D0F2xF4_x\[FF:00\]](#). To read or write to one of these register, the address is written first into the address register [D0F2xF0](#) and then the data are read or written by read or write the data register [D0F2xF4](#). See [2.12.1 \[IOMMU Configuration Space\]](#).

Bits	Description
31:9	Reserved.
8	L2cfgWrEn . Read-write. Reset: 0. 1=Enable writes to D0F2xF4 .
7:0	L2cfgIndex . Read-write. Reset: 0.

D0F2xF4 IOMMU L2 Config DataSee [D0F2xF0](#).**D0F2xF4_x00 L2_PERF_CNTL_0**

Bits	Description
31:24	L2PerfCountUpper1 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 1
23:16	L2PerfCountUpper0 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 0
15:8	L2PerfEvent1 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 1
7:0	L2PerfEvent0 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 0

D0F2xF4_x01 L2_PERF_COUNT_0

Bits	Description
31:0	L2PerfCount0 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 0

D0F2xF4_x02 L2_PERF_COUNT_1

Bits	Description
31:0	L2PerfCount1 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 1

D0F2xF4_x03 L2_PERF_CNTL_1

Bits	Description
31:24	L2PerfCountUpper3 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 3
23:16	L2PerfCountUpper2 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 2
15:8	L2PerfEvent3 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 3
7:0	L2PerfEvent2 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 2

D0F2xF4_x04 L2_PERF_COUNT_2

Bits	Description
31:0	L2PerfCount2 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 2

D0F2xF4_x05 L2_PERF_COUNT_3

Bits	Description
31:0	L2PerfCount3 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 3

D0F2xF4_x06 L2_DEBUG_0

Bits	Description
31:0	L2DEBUG0 . Read-write. Reset: 0. Reserved for to control ECOs

D0F2xF4_x07 L2_DEBUG_1

Bits	Description
31:0	L2DEBUG1 . Read-write. Reset: 0. Reserved for to control ECOs

D0F2xF4_x08 L2_STATUS_0

Bits	Description
31:0	L2STATUS0 . Read-only. Reset: 0. Internal IOMMU L2A status

D0F2xF4_x0C L2_CONTROL_0

Bits	Description
31:24	IFifoClientPriority . Read-write. Reset: 0. Each bit of this register controls whether the corresponding L1 client is arbitrated as high priority or not. Not all implementations will use all of the priority bits due to a lower number of clients versus the register width
23:20	IFifoBurstLength . Read-write. Reset: 1. Sets the burst length when arbitrating between clients coming into the L2
19	Reserved.
18	FLTCMBPriority . Read-write. Reset: 0. 0=Round-robin arbitration between cache responses and table-walker responses at the fault combiner. . 1=Table-walker responses always win arbitration at the fault combiner.
17:12	IFifoCMBCredits . Read-write. Reset: 4h. Controls the initial number of credits for the ififo to fault/CMB interface. Credits are loaded whenever the register value changes. This register may only be programmed when IOMMU is not enabled to preserve correct operation.
11	SIDEPTEOnAddrTransExcl . Read-write. Reset: 0. 0=Caches return DTE to L1 on an address translation exclusion range access. 1=Caches return PTE to L1 on an address translation exclusion range access
10	SIDEPTEOnUntransExcl . Read-write. Reset: 0. 0=Caches return DTE to L1 on an untranslated exclusion range access . 1=Caches return PTE to L1 on an untranslated exclusion range access

9:4	IFifoTWCredits. Read-write. Reset: 4h. Controls the initial number of credits for the ififo to TW interface. Credits are loaded whenever the register changes value. This register may only be programmed when IOMMU is not enabled to preserve correct operation.
3	DTCHitVZeroOrIVZero. Read-write. Reset: 0. 0=A DTE is refetched if a DTE with V=0 for a memory request or IV=1 for an interrupt request is hit in the DTC . 1=A DTE is not refetched if a DTE with V=0 for a memory request or IV=1 for an interrupt request is hit in the DTC. This DTE is used
2	AllowL1CacheATSRsp. Read-write. Reset: 0. 0=L2 does not allow L1 to cache responses to ATS address translation requests . 1=L2 allows L1 to cache responses to ATS address translation requests
1	AllowL1CacheVZero. Read-write. Reset: 0. 0=L2 does not allow L1 to cache DTEs where V=0 . 1=L2 allows L1 to cache DTEs where V=1. L1 stores IR and IW as if they are both set to 1
0	PTCAddrTransReqCheck. Read-write. Reset: 0. 0=Address translation requests do not check the PTC . 1=Address translation requests check the PTC

D0F2xF4_x0D L2_CONTROL_1

Bits	Description
31:24	PerfThreshold. Read-write. Reset: 0. Fifo threshold level used to calculate certain performance counter values.
23:17	Reserved.
16	SeqInvBurstLimitEn. Read-write. Reset: 1. Enable stalling L2 requests to allow invalidation cycles to make forward progress based upon SeqInvBurstLimitInv and SeqInvBurstLimitL2Req
15:8	SeqInvBurstLimitL2Req. Read-write. Reset: 8h. Sets the number of consecutive IOMMU L2 requests to perform when doing sequential invalidation. Regular L2 and invalidation requests will alternate access to the main L2 caches based upon SeqInvBurstLimitInv and SeqInvBurstLimitL2Req.
7:0	SeqInvBurstLimitInv. Read-write. Reset: 8h. Sets the number of consecutive invalidation requests to perform when doing sequential invalidation. Regular L2 and invalidation requests will alternate access to the main L2 caches based upon SeqInvBurstLimitInv and SeqInvBurstLimitL2Req.

D0F2xF4_x10 L2_DTC_CONTROL

Bits	Description
31:28	DTCEntries. Read-only. Reset: 0. The number of entries in the DTC is indicated as $2^{\text{DTCEntries}}$
27:24	Reserved.
23:16	DTCWays. Read-only. Reset: 0. Indicates the number of ways in the DTC
15	DTCParitySupport. Read-only. Reset: 0. 0=The DTC does not support parity protection . 1=The DTC supports parity protection
14	Reserved.
13	DTCBypass. Read-write. Reset: 0. When set, all requests bypass the DTC.
12:11	Reserved.

10	DTCSoftInvalidate. Read-write. Reset: 0. Software may write this register to 1 to invalidate all entries in the DTC										
9:8	DTCInvalidationSel. Read-write. Reset: 0. BIOS: 10b. Selects the DTC invalidation algorithm. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Invalidate the entire DTC.</td> </tr> <tr> <td>01b</td> <td>Fast imprecise invalidation.</td> </tr> <tr> <td>10b</td> <td>Sequential precise invalidation.</td> </tr> <tr> <td>11b</td> <td>Partial sequential precise invalidation.</td> </tr> </tbody> </table>	Bits	Description	00b	Invalidate the entire DTC.	01b	Fast imprecise invalidation.	10b	Sequential precise invalidation.	11b	Partial sequential precise invalidation.
Bits	Description										
00b	Invalidate the entire DTC.										
01b	Fast imprecise invalidation.										
10b	Sequential precise invalidation.										
11b	Partial sequential precise invalidation.										
7:5	Reserved.										
4	DTCParityEn. Read-write. Reset: 0. Enable parity protection of the DTC										
3	DTCLRUUpdatePri. Read-write. Reset: 0. 0=Reads update replacement state bits when there is a simultaneous read and write to the same DTC index . 1=Writes update replacement state bits when there is a simultaneous read and write to the same DTC index										
2	Reserved.										
1:0	DTCReplacementSel. Read-write. Reset: 1. Selects the DTC replacement algorithm. Implementation may not support all replacement algorithms										

D0F2xF4_x11 L2_DTC_HASH_CONTROL

Bits	Description
31:16	DtcAddressMask. Read-write. Reset: FFFFh. BIOS: 0. This field is a bit-wise AND mask that selects which bits from the untranslated interrupt {MT[2:0],Vector} are used to index into the DTC.
15:11	Reserved.
10	DtcAltHashEn. Read-write. Reset: 0. BIOS: 1. Enable alternative algorithm for generating hash index into the DTC.
9	Reserved.
8:5	DTCBusBits. Read-write. Reset: 3h. Set the number of bus bits to use when using ReqID to form the DTC address. The following equation must be satisfied. $\text{Func_bits} + \text{Dev_Bits} + \text{Bus_Bits} \leq \log_2(\text{DTC entries} / \text{DTC associativity})$.
4:2	DTCDevBits. Read-write. Reset: 0. Set the number of device bits to use when using ReqID to form the DTC address.
1:0	DTCFuncBits. Read-write. Reset: 2h. Set the number of function bits to use when using ReqID to form the DTC address.

D0F2xF4_x12 L2_DTC_WAY_CONTROL

Bits	Description
31:16	DTCWayAccessDisable. Read-write. Reset: 0.
15:0	DTCWayDisable. Read-write. Reset: 0. Each bit in this register disables a way in the DTC when set to 1. An implementation may have less than 32 ways. The entire cache may be disabled by setting the DTCWays lower bits of this register.

D0F2xF4_x14 L2_ITC_CONTROL

Bits	Description										
31:28	ITCEntries . Read-only. Reset: 0. The number of entries in the ITC is indicated as $2^{\text{ITCEntries}}$										
27:24	Reserved.										
23:16	ITCWays . Read-only. Reset: 0. Indicates the number of ways in the ITC										
15	ITCParitySupport . Read-only. Reset: 0. 0=The ITC does not support parity protection . 1=The ITC supports parity protection										
14	Reserved.										
13	ITCBypass . Read-write. Reset: 0. When set, all requests bypass the ITC.										
12:11	Reserved.										
10	ITCSoftInvalidate . Read-write. Reset: 0. Software may write this register to 1 to invalidate all entries in the ITC										
9:8	ITCInvalidationSel . Read-write. Reset: 0. BIOS: 10b. Selects the ITC invalidation algorithm. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Invalidate the entire ITC.</td> </tr> <tr> <td>01b</td> <td>Fast imprecise invalidation.</td> </tr> <tr> <td>10b</td> <td>Sequential precise invalidation.</td> </tr> <tr> <td>11b</td> <td>Partial sequential precise invalidation</td> </tr> </tbody> </table>	Bits	Description	00b	Invalidate the entire ITC.	01b	Fast imprecise invalidation.	10b	Sequential precise invalidation.	11b	Partial sequential precise invalidation
Bits	Description										
00b	Invalidate the entire ITC.										
01b	Fast imprecise invalidation.										
10b	Sequential precise invalidation.										
11b	Partial sequential precise invalidation										
7:5	Reserved.										
4	ITCParityEn . Read-write. Reset: 0. Enable parity protection of the ITC										
3	ITCLRUpdatePri . Read-write. Reset: 0. 0=Reads update replacement state bits when there is a simultaneous read and write to the same ITC index . 1=Writes update replacement state bits when there is a simultaneous read and write to the same ITC index										
2	Reserved.										
1:0	ITCReplacementSel . Read-write. Reset: 1. Selects the ITC replacement algorithm. Implementation may not support all replacement algorithms										

D0F2xF4_x15 L2_ITC_HASH_CONTROL

Bits	Description
31:16	ITCAddressMask . Read-write. Reset: FFFFh. BIOS: 0. This register is a bit-wise AND mask that selects which bits from the untranslated interrupt {MT[2:0],Vector} are used to index into the ITC.
15:11	Reserved.
10	ItcAltHashEn . Read-write. Reset: 0. BIOS: 1. Enable alternative algorithm for generating hash index into the ITC.
9	Reserved.
8:5	ITCBusBits . Read-write. Reset: 3h. Set the number of bus bits to use when using ReqID to form the ITC address. The following equation must be satisfied. $\text{Func_bits} + \text{Dev_Bits} + \text{Bus_Bits} \leq \log_2(\text{ITC entries} / \text{ITC associativity})$

4:2	ITCDevBits . Read-write. Reset: 0. Set the number of device bits to use when using ReqID to form the ITC address.
1:0	ITCFuncBits . Read-write. Reset: 2h. Set the number of function bits to use when using ReqID to form the ITC address.

D0F2xF4_x16 L2_ITC_WAY_CONTROL

Bits	Description
31:16	ITCWayAccessDisable . Read-write. Reset: 0.
15:0	ITCWayDisable . Read-write. Reset: 0. Each bit in this register disables a way in the ITC when set to 1. An implementation may have less than 32 ways. The entire cache may be disabled by setting the ITCWays lower bits of this register.

D0F2xF4_x18 L2_PTC_A_CONTROL

Bits	Description										
31:28	PTCAEntries . Read-only. Reset: 0. The number of entries in the PTC A sub-cache is indicated as $2^{\text{PTCAEntries}}$										
27:24	Reserved.										
23:16	PTCAWays . Read-only. Reset: 0. Indicates the number of ways in the PTC A sub-cache										
15	PTCAParitySupport . Read-only. Reset: 0. 0=The PTC A sub-cache does not support parity protection . 1=The PTC A sub-cache supports parity protection										
14	Reserved.										
13	PTCABypass . Read-write. Reset: 0. When set, all requests bypass the PTC A sub-cache.										
12	Reserved.										
11	PTCA2MMode . Read-write. Reset: 0. When set, the PTC A sub-cache stores 2M pages instead of 4K pages										
10	PTCASoftInvalidate . Read-write. Reset: 0. Software may write this register to 1 to invalidate all entries in the PTC A sub-cache										
9:8	PTCAInvalidationSel . Read-write. Reset: 0. BIOS: 10b. Selects the PTC A sub-cache invalidation algorithm. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Invalidate the entire PTC A sub-cache.</td> </tr> <tr> <td>01b</td> <td>Fast imprecise invalidation.</td> </tr> <tr> <td>10b</td> <td>Sequential precise invalidation.</td> </tr> <tr> <td>11b</td> <td>Partial sequential precise invalidation.</td> </tr> </tbody> </table>	Bits	Description	00b	Invalidate the entire PTC A sub-cache.	01b	Fast imprecise invalidation.	10b	Sequential precise invalidation.	11b	Partial sequential precise invalidation.
Bits	Description										
00b	Invalidate the entire PTC A sub-cache.										
01b	Fast imprecise invalidation.										
10b	Sequential precise invalidation.										
11b	Partial sequential precise invalidation.										
7:5	Reserved.										
4	PTCAParityEn . Read-write. Reset: 0. Enable parity protection of the PTC A sub-cache										
3	PTCALRUUpdatePri . Read-write. Reset: 0. 0=Reads update replacement state bits when there is a simultaneous read and write to the same PTCA index . 1=Writes update replacement state bits when there is a simultaneous read and write to the same PTCA index										

2	Reserved.
1:0	PTCAReplacementSel. Read-write. Reset: 1. Selects the PTC A sub-cache replacement algorithm. Implementation may not support all replacement algorithms

D0F2xF4_x19 L2_PTC_A_HASH_CONTROL

Bits	Description
31:16	PTCAAddressMask. Read-write. Reset: FFFFh. BIOS: 0. This register is a bit-wise AND mask that selects which virtual address bits are used to index into the PTC A sub-cache.
15:11	Reserved.
10	PtcAltHashEn. Read-write. Reset: 0. BIOS: 1. Enable alternative algorithm for generating hash index into the PTC
9	Reserved.
8:5	PTCABusBits. Read-write. Reset: 3h. Set the number of bus bits to use when using ReqID to form the PTC A sub-cache address. The following equation must be satisfied. $\text{FuncBits} + \text{DevBits} + \text{BusBits} \leq \log_2(\text{PTC A sub-cache entries} / \text{PTC A sub-cache associativity})$
4:2	PTCADevBits. Read-write. Reset: 0. Set the number of device bits to use when using ReqID to form the PTC A sub-cache address
1:0	PTCAFuncBits. Read-write. Reset: 2h. Set the number of function bits to use when using ReqID to form the PTC A sub-cache address

D0F2xF4_x1A L2_PTC_A_WAY_CONTROL

Bits	Description
31:16	PTCAWayAccessDisable. Read-write. Reset: 0.
15:0	PTCAWayDisable. Read-write. Reset: 0. Each bit in this register disables a way in the PTC A sub-cache when set to 1. An implementation may have less than 32 ways. The entire cache may be disabled by setting the PTCAWays lower bits of this register.

D0F2xF4_x20 L2_CREDIT_CONTROL_2

Bits	Description
31:28	Reserved.
27:24	PprLoggerCredits. Read-write. Reset: 4h. PPR log buffer credit override value
23	FCELOverride. Read-write. Reset: 0. Changing this register from 0 to 1 overrides the FCEL credit counter with FCELCredits. This should only be performed when the IOMMU is idle
22	Reserved.
21:16	FCELCredits. Read-write. Reset: 0. FCEL credit override value
15	FLTCMBOverride. Read-write. Reset: 0. Changing this register from 0 to 1 overrides the FLTCMB credit counter with FLTCMBCredits. This should only be performed when the IOMMU is idle
14	Reserved.

13:8	FLTCMBCredits . Read-write. Reset: 0. FLTCMB credit override value
7	QUEUEOverride . Read-write. Reset: 0. Changing this register from 0 to 1 overrides the QUEUE credit counter with QUEUECredits. This should only be performed when the IOMMU is idle
6	Reserved.
5:0	QUEUECredits . Read-write. Reset: 0. QUEUE credit override value

D0F2xF4_x22 L2A_UPDATE_FILTER_CNTL

Bits	Description
31:5	Reserved.
4:1	L2aUpdateFilterRdlatency . Read-write. Reset: 3h. When L2a_Update_Filter_Bypass is 0, assume the invalidation read has completed in the number of clock cycles specified by this field.
0	L2aUpdateFilterBypass . Read-write. Reset: 1. 1=Disable duplicate update filtering. 0=Enable the dropping of updates that are already in the L2aUpdateFilter or in the destination L2a cache.

D0F2xF4_x30 L2_ERR_RULE_CONTROL_3

Bits	Description
31:4	ERRRuleDisable3 . Read-write. Reset: 0. Each bit in this register disables an error detection rule in the IOMMU.
3:1	Reserved.
0	ERRRuleLock . Read-write. Reset: 0. BIOS: 1. This register is write-once. Setting this register bit locks the error detection rule set in ERRRuleDisable3 , D0F2xF4_x31[ERRRuleDisable4] , and D0F2xF4_x32[ERRRuleDisable5] .

D0F2xF4_x31 L2_ERR_RULE_CONTROL_4

Bits	Description
31:0	ERRRuleDisable4 . Read-write. Reset: 0. Each bit in this register disables an error detection rule in the IOMMU.

D0F2xF4_x32 L2_ERR_RULE_CONTROL_5

Bits	Description
31:0	ERRRuleDisable5 . Read-write. Reset: 0. Each bit in this register disables an error detection rule in the IOMMU.

D0F2xF4_x33 L2_L2A_CK_GATE_CONTROL

Bits	Description
31:8	Reserved.

7:6	CKGateL2AStop . Read-write. Reset: 01b. 0=allow 2 clock cycles delay before stopping the clocks when clkready deasserts . 1=allow 4 clock cycles delay before stopping the clocks when clkready deasserts. 2=allow 8 clock cycles delay before stopping the clocks when clkready deasserts . 3=allow 16 clock cycles delay before stopping the clocks when clkready deasserts.
5:4	CKGateL2ALength . Read-write. Reset: 01b. 0=allow 128 clock cycles delay before stopping the clocks when idle asserts . 1=allow 256 clock cycles delay before stopping the clocks when idle asserts. 2=allow 512 clock cycles delay before stopping the clocks when idle asserts . 3=allow 1024 clock cycles delay before stopping the clocks when idle asserts.
3	CKGateL2ASpare . Read-write. Reset: 0.
2	CKGateL2ACacheDisable . Read-write. Reset: 1. BIOS: 0. Disable the gating of the l2b upper cache ways.
1	CKGateL2ADynamicDisable . Read-write. Reset: 1. BIOS: 0. Disable the gating of the l2b dynamic clock branch.
0	CKGateL2ARegsDisable . Read-write. Reset: 1. BIOS: 0. Disable the gating of the l2b register clock branch.

D0F2xF4_x34 L2_L2A_PGFSIZE_CONTROL

Bits	Description
31:4	Reserved.
3:2	L2aregHostPgsize . Read-write. Reset: 0. BIOS: 10b.
1:0	L2aregGstPgsize . Read-write. Reset: 0. BIOS: 10b.

D0F2xF4_x40 L2_PERF_CNTL_2

Bits	Description
31:24	L2PerfCountUpper5 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 5
23:16	L2PerfCountUpper4 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 4
15:8	L2PerfEvent5 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 5
7:0	L2PerfEvent4 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 4

D0F2xF4_x41 L2_PERF_COUNT_4

Bits	Description
31:0	L2PerfCount4 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 4

D0F2xF4_x42 L2_PERF_COUNT_5

Bits	Description
31:0	L2PerfCount5 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 5

D0F2xF4_x43 L2_PERF_CNTL_3

Bits	Description
31:24	L2PerfCountUpper7 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 7
23:16	L2PerfCountUpper6 . Read-only. Reset: 0. Upper 8 bits of IOMMU L2 performance counter 6
15:8	L2PerfEvent7 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 7
7:0	L2PerfEvent6 . Read-write. Reset: 0. Selects the IOMMU L2 performance counter event for counter 6

D0F2xF4_x44 L2_PERF_COUNT_6

Bits	Description
31:0	L2PerfCount6 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 6

D0F2xF4_x45 L2_PERF_COUNT_7

Bits	Description
31:0	L2PerfCount7 . Read-only. Reset: 0. Lower 32 bits of IOMMU L2 performance counter 7

D0F2xF4_x46 L2_DEBUG_2

Bits	Description
31:0	L2DEBUG2 . Read-write. Reset: 0. Reserved for to control ECOs

D0F2xF4_x47 L2_DEBUG_3

Bits	Description
31:3	Reserved.
2	TwAtomicFilterEn . Read-write. Reset: 0. BIOS: 1. 1=Enable table walker atomic filtering.
1	TwNwEn . Read-write. Reset: 0. BIOS: 1. 1=Enable NW bit for ATS requests.
0	Reserved.

D0F2xF4_x48 L2_STATUS_1

Bits	Description
31:0	L2STATUS1 . Read-only. Reset: 0. Internal IOMMU L2B status

D0F2xF4_x4C L2_CONTROL_5

Bits	Description
31:10	Reserved. Read-write.
9:8	GstPartialPtcCntl. Read-write. Reset: 0. BIOS: 11b.
7	FC2AltMode. Read-write. Reset: 0. 0=FC2 primary flow-control mode . 1=FC2 alternate flow-control mode
6	FC3Dis. Read-write. Reset: 0. 0=FC3 flow-control loop is enabled . 1=FC3 flow-control look is disabled
5	FC2Dis. Read-write. Reset: 0. 0=FC2 flow-control loop is enabled . 1=FC2 flow-control look is disabled
4	DTCUpdateVZeroIVOne. Read-write. Reset: 0. 0=DTEs with V=0 and IV=1 are not cached in the DTC . 1=DTEs with V=0 and IV=1 are cached in the DTC
3	DTCUpdateVOneIVZero. Read-write. Reset: 0. 0=DTEs with V=1 and IV=0 are not cached in the DTC . 1=DTEs with V=1 and IV=0 are cached in the DTC
2	FC1Dis. Read-write. Reset: 0. 0=FC1 flow control loop enabled . 1=FC1 flow control loop disabled
1	PTCAddrTransReqUpdate. Read-write. Reset: 1. 0=PTEs from address translation requests are not cached . 1=PTEs from address translation requests are cached in the L2 according to the Cache bit in the DTE
0	QueueArbFBPri. Read-write. Reset: 1. 0=Requests in the miss queue and the feedback queue are arbitrated in a round-robin manner . 1=Requests in the feedback queue are given priority over requests in the miss queue

D0F2xF4_x4D L2_CONTROL_6

Bits	Description
31:24	Perf2Threshold. Read-write. Reset: 0. Fifo threshold level used to calculate certain performance counter values.
23:17	Reserved.
16	SeqInvBurstLimitEn. Read-write. Reset: 1. Enable stalling PDC requests to allow invalidation cycles to make forward progress based upon SeqInvBurstLimitInv and SeqInvBurstLimitPDCReq
15:8	SeqInvBurstLimitPDCReq. Read-write. Reset: 8h. Sets the number of consecutive IOMMU PDC requests to perform when doing sequential invalidation. PDC and invalidation requests will alternate access to the PDC based upon SeqInvBurstLimitInv and SeqInvBurstLimitPDCReq.
7:0	SeqInvBurstLimitInv. Read-write. Reset: 8h. Sets the number of consecutive invalidation requests to perform when doing sequential invalidation. PDC and invalidation requests will alternate access to the PDC based upon SeqInvBurstLimitInv and SeqInvBurstLimitPDCReq.

D0F2xF4_x50 L2_PDC_CONTROL

Bits	Description										
31:28	PDCEntries . Read-only. Reset: 0. Indicates the number of entries in the PDC is indicated as $2^{\text{PDCEntries}}$										
27:24	Reserved.										
23:16	PDCWays . Read-only. Reset: 0. Indicates the number of ways in the PDC										
15	PDCParitySupport . Read-only. Reset: 0. 0=The PDC does not support parity protection . 1=The PDC supports parity protection										
14	Reserved.										
13	PDCBypass . Read-write. Reset: 0. When set, all requests bypass the PDC. This prevents the multiple issue of requests and increases maximum rate of requests to the table-walker.										
12	PDCSearchDirection . Read-write. Reset: 0. 0=Search PDC from higher levels down . 1=Search PDC from lower levels up										
11	Reserved.										
10	PDCSoftInvalidate . Read-write. Reset: 0. Software may write this register to 1 to invalidate all entries in the PDC										
9:8	PDCInvalidationSel . Read-write. Reset: 0. BIOS: 10b. Selects the PDC invalidation algorithm. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Invalidate the entire PDC.</td> </tr> <tr> <td>01b</td> <td>Fast imprecise invalidation.</td> </tr> <tr> <td>10b</td> <td>Sequential precise invalidation.</td> </tr> <tr> <td>11b</td> <td>Partial sequential precise invalidation</td> </tr> </tbody> </table>	Bits	Description	00b	Invalidate the entire PDC.	01b	Fast imprecise invalidation.	10b	Sequential precise invalidation.	11b	Partial sequential precise invalidation
Bits	Description										
00b	Invalidate the entire PDC.										
01b	Fast imprecise invalidation.										
10b	Sequential precise invalidation.										
11b	Partial sequential precise invalidation										
7:5	Reserved.										
4	PDCParityEn . Read-write. Reset: 0. Enable parity protection of the PDC if the device supports parity										
3	PDCLRUUpdatePri . Read-write. Reset: 0. 0=Reads update replacement state bits when there is a simultaneous read and write to the same PDC index . 1=Writes update replacement state bits when there is a simultaneous read and write to the same PDC index										
2	Reserved.										
1:0	PDCReplacementSel . Read-write. Reset: 1. Selects the PDC replacement algorithm. Implementation may not support all replacement algorithms										

D0F2xF4_x51 L2_PDC_HASH_CONTROL

Bits	Description
31:16	PDCAddressMask . Read-write. Reset: FFFFh. BIOS: 0. This register is a bit-wise AND mask that selects which virtual address bits are used to index into the PDC.
15:11	Reserved.
10	PdcAltHashEn . Read-write. Reset: 0. BIOS: 1. 1=Enable alternative algorithm for generating hash index into the PDC.

9	PDCUpperLvlAddrHash. Read-write. Reset: 1. When set to 1, the PDC cache index is partially formed using the xor of the LSBs of virtual address bits for all levels greater than or equal to the stored/searched level.
8	PDCLvlHash. Read-write. Reset: 1. When set to 1, the PDE level is used as part of the hash for the cache index.
7:6	Reserved.
5:0	PDCDomainBits. Read-write. Reset: 7h. Selects the number of domain bits to use as part of the index into the PDC.

D0F2xF4_x52 L2_PDC_WAY_CONTROL

Bits	Description
31:16	PDCWayAccessDisable. Read-write. Reset: 0.
15:0	PDCWayDisable. Read-write. Reset: 0. Each bit in this register disables a way in the PDC when set to 1. An implementation may have less than 32 ways. The entire cache may be disabled by setting the PDCWays lower bits of this register.

D0F2xF4_x53 L2B_UPDATE_FILTER_CNTL

Bits	Description
31:5	Reserved.
4:1	L2bUpdateFilterRdlatency. Read-write. Reset: 3h. When L2b_Update_Filter_Bypass is 0, assume the invalidation read has completed in the number of clock cycles specified by this field.
0	L2bUpdateFilterBypass. Read-write. Reset: 1. 1 - Disable duplicate update filtering; 0 - Enable the dropping of updates that are already in the l2b_update_filter or in the PDC.

D0F2xF4_x54 L2_TW_CONTROL

Bits	Description
31:18	Reserved.
17	Twfilter64bDis. Read-write. Reset: 0.
16	TwfilterDis. Read-write. Reset: 0.
15	Reserved.
14:12	TWPrefetchRange. Read-write. Reset: 1. Selects the number of pages to prefetch
11	TWPTEOnAddrTransExcl. Read-write. Reset: 0. 0=Table walker returns DTE to L1 on an address translation exclusion range access . 1=Table walker returns PTE to L1 on an address translation exclusion range access
10	TWPTEOnUntransExcl. Read-write. Reset: 0. 0=Table walker returns DTE to L1 on an untranslated exclusion range access . 1=Table walker returns PTE to L1 on an untranslated exclusion range access
9	TWPrefetchOnly4KDis. Read-write. Reset: 0. 1=Allow non-4K pages to be prefetched . 0=Only 4K pages are prefetched

8	TWPrefetchEn. Read-write. Reset: 0. Enable prefetching in the table-walker
7	Reserved.
6	TWForceCoherent. Read-write. Reset: 0. 1=Table-walker always generates coherent requests. The DTE SD bit is ignored when this bit is set to 1.
5:0	TWQueueLimit. Read-write. Reset: 10h. Limit the number of outstanding table-walker requests

D0F2xF4_x56 L2_CP_CONTROL

Bits	Description
31:16	CPRdDelay. Read-write. Reset: 0. Command processor read delay
15:3	Reserved.
2	CPFlushOnInv. Read-write. Reset: 1. BIOS: 0. 1=Command processor flushes out old requests on every invalidation command . 0=No flush is performed during invalidations
1	CPFlushOnWait. Read-write. Reset: 0. BIOS: 1. 1=Command processor flushes out old requests on completion wait . 0=No flush is performed on completion wait
0	CPPrefetchDis. Read-write. Reset: 0. 1=Command processor fetches and executes only one command at a time . 0=Command processor prefetches available commands into its internal storage

D0F2xF4_x57 L2_CP_CONTROL_1

Bits	Description
31:3	Reserved.
2	L1ImuIntGfxDis. Read-write. Reset: 0. BIOS: IF (GpuEnabled) THEN 0 ELSE 1 ENDIF.
1	Reserved.
0	L1ImuPcieGfxDis. Read-write. Reset: 0. BIOS: This bit should be set if there is no external graphics in the system.

D0F2xF4_x58 IOMMU_L2_GUEST_ADDR_CNTRL

Bits	Description
31:24	Reserved.
23:0	IommuL2GuestAddrMask. Read-write. Reset: 0.

D0F2xF4_x60 L2_TW_CONTROL_1

Bits	Description
31:15	TWDebugMask. Read-write. Reset: 0. Defines the table-walker trace buffer size by masking address bits 31:16
14:3	Reserved.

2	TWDebugForceDisable. Read-write. Reset: 0. When set to 1, this register disables the TW Debug feature until a cold-reset is performed
1	TWDebugNoWrap. Read-write. Reset: 0. 1=Table-walker trace to stop at the top of the trace buffer . 0=Table-walker trace wraps around at the top of the trace buffer
0	TWDebugEn. Read-write. Reset: 0. Enables table-walker trace debug mode

D0F2xF4_x61 L2_TW_CONTROL_2

Bits	Description
31:12	TWDebugAddrLo. Read-write. Reset: 0. Base address bits 31:12 for table-walker trace debug
11:0	Reserved.

D0F2xF4_x62 L2_TW_CONTROL_3

Bits	Description
31:0	TWDebugAddrHi. Read-write. Reset: 0. Base address bits 51:32 for table-walker trace debug

D0F2xF4_x6A L2_INT_CONTROL

Bits	Description
31:3	Reserved.
2	IntPPROrderEn. Read-write. Reset: 1. Enable ordering between interrupts and ppr log writes
1	IntCPOrderEn. Read-write. Reset: 1. Enable ordering between interrupts and command processor writes
0	IntEventOrderEn. Read-write. Reset: 1. Enable ordering between interrupts and event log writes

D0F2xF4_x70 L2_CREDIT_CONTROL_0

Bits	Description
31	DTEOverride. Read-write. Reset: 0. Changing this register from 0 to 1 overrides the DTE credit counter with DTECredits. This should only be performed when the IOMMU is idle
30	Reserved.
29:24	DTECredits. Read-write. Reset: 0. DTE credit override value
23	FC3Override. Read-write. Reset: 0. Changing this register from 0 to 1 overrides the FC3 credit counter with FC3Credits. This should only be performed when the IOMMU is idle
22	Reserved.
21:16	FC3Credits. Read-write. Reset: 0. FC3 credit override value
15	FC2Override. Read-write. Reset: 0. Changing this register from 0 to 1 overrides the FC2 credit counter with FC2Credits. This should only be performed when the IOMMU is idle
14	Reserved.
13:8	FC2Credits. Read-write. Reset: 0. FC2 credit override value

7	FC1Override . Read-write. Reset: 0. Changing this register from 0 to 1 overrides the FC1 credit counter with FC1Credits. This should only be performed when the IOMMU is idle
6	Reserved.
5:0	FC1Credits . Read-write. Reset: 0. FC1 credit override value

D0F2xF4_x71 L2_CREDIT_CONTROL_1

Bits	Description
31:24	Reserved.
23:20	PprMcifCredits . Read-write. Reset: 4h. PPR logger credit override value
19:16	CpPrefetchCredits . Read-write. Reset: 4h. Command processor prefetch credit override value
15	TWELOverride . Read-write. Reset: 0. Changing this register from 0 to 1 overrides the TWEL credit counter with TWELCredits. This should only be performed when the IOMMU is idle
14	Reserved.
13:8	TWELCredits . Read-write. Reset: 0. TWEL credit override value
7	PDTIEOverride . Read-write. Reset: 0. Changing this register from 0 to 1 overrides the PDTIE credit counter with PDTIECredits. This should only be performed when the IOMMU is idle
6	Reserved.
5:0	PDTIECredits . Read-write. Reset: 0. PDTIE credit override value

D0F2xF4_x78 L2_MCIF_CONTROL

Bits	Description
31:29	Reserved.
28:24	MCIFBaseWriteDataCredits . Read-write. Reset: 8h. Sets the number of base-channel write data credits between the IOMMU L2 and the HTIU/ORB. This register requires a warm-reset to take effect
23:21	Reserved.
20:16	MCIFBaseWriteHdrCredits . Read-write. Reset: 8h. Sets the number of base-channel write header credits between the IOMMU L2 and the HTIU/ORB. This register requires a warm-reset to take effect
15:13	Reserved.
12:8	MCIFIsocReadCredits . Read-write. Reset: 8h. Sets the number of isoc-channel read credits between the IOMMU L2 and the HTIU/ORB. This register requires a warm-reset to take effect
7:5	Reserved.
4:0	MCIFBaseReadCredits . Read-write. Reset: 8h. Sets the number of base-channel read credits between the IOMMU L2 and the HTIU/ORB. This register requires a warm-reset to take effect

D0F2xF4_x80 L2_ERR_RULE_CONTROL_0

Bits	Description
31:4	ERRRuleDisable0 . Read-write. Reset: 0. Each bit in this register disables an error detection rule in the IOMMU.
3:1	Reserved.
0	ERRRuleLock0 . Read-write. Reset: 0. BIOS: 1. This register is write-once. Setting this register bit locks the error detection rule set in ERRRuleDisable0 , D0F2xF4_x81 [ERRRuleDisable1], and D0F2xF4_x82 [ERRRuleDisable2].

D0F2xF4_x81 L2_ERR_RULE_CONTROL_1

Bits	Description
31:0	ERRRuleDisable1 . Read-write. Reset: 0. Each bit in this register disables an error detection rule in the IOMMU.

D0F2xF4_x82 L2_ERR_RULE_CONTROL_2

Bits	Description
31:0	ERRRuleDisable2 . Read-write. Reset: 0. Each bit in this register disables an error detection rule in the IOMMU.

D0F2xF4_x90 L2_L2B_CK_GATE_CONTROL

Bits	Description
31:8	Reserved.
7:6	CKGateL2BStop . Read-write. Reset: 01b. 0=allow 2 clock cycles delay before stopping the clocks when clkready deasserts . 1=allow 4 clock cycles delay before stopping the clocks when clkready deasserts . 2=allow 8 clock cycles delay before stopping the clocks when clkready deasserts . 3=allow 16 clock cycles delay before stopping the clocks when clkready deasserts
5:4	CKGateL2BLength . Read-write. Reset: 01b. 0=allow 128 clock cycles delay before stopping the clocks when idle asserts . 1=allow 256 clock cycles delay before stopping the clocks when idle asserts . 2=allow 512 clock cycles delay before stopping the clocks when idle asserts . 3=allow 1024 clock cycles delay before stopping the clocks when idle asserts
3	CKGateL2BCacheDisable . Read-write. Reset: 0. Disable the gating of the l2b upper cache ways
2	CKGateL2BMiscDisable . Read-write. Reset: 1. Disable the gating of the l2b miscellaneous clock branch
1	CKGateL2BDynamicDisable . Read-write. Reset: 1. BIOS: 0. Disable the gating of the l2b dynamic clock branch
0	CKGateL2BRegsDisable . Read-write. Reset: 1. BIOS: 0. Disable the gating of the l2b register clock branch

D0F2xF4_x92 PPR_CONTROL

Bits	Description
31:17	Reserved.
16	PprIntcoalesceEn. Read-write. Reset: 0. BIOS: 0. This bit must not be set if D0F2xF4_x90[CKGateL2BMiscDisable]=0 .
15:8	PprIntreqdelay. Read-write. Reset: 0. BIOS: 20h.
7:0	PprInttimedelay. Read-write. Reset: 0. BIOS: 15h.

D0F2xF4_x94 L2_L2B_PGSIZE_CONTROL

Bits	Description
31:4	Reserved.
3:2	L2bregHostPgsz. Read-write. Reset: 0. BIOS: 10b.
1:0	L2bregGstPgsz. Read-write. Reset: 0. BIOS: 10b.

D0F2xF8 IOMMU L1 Config Index

The index/data pair registers [D0F2xF8](#) and [D0F2xFC](#) is used to access the registers [D0F2xFC_x\[FFFF:0000\]_L1sel\[3:0\]](#). To read or write to one of these register, the address is written first into the address register [D0F2xF8](#) and then the data are read or written by read or write the data register [D0F2xFC](#).

See [2.12.1 \[IOMMU Configuration Space\]](#). There are four L1s in the IOMMU. Registers in the L1 indexed space have one instance per L1 denoted by [_L1sel\[x\]](#) where $x = \text{D0F2xF8[L1cfgSel]}$. The syntax for this register type is described by example as follows:

- [D0F2xFC_x00_L1sel\[3:0\]](#) refers to all instances of the [D0F2xFC_x00](#) registers.
- [D0F2xFC_x00_L1sel\[2\]](#) refers to the GBIF instance of [D0F2xFC_x00](#).

Bits	Description												
31	L1cfgEn. Read-write. Reset: 0. 1=Enable writes to D0F2xFC .												
30:20	Reserved.												
19:16	L1cfgSel. Read-write. Reset: 0. This field selects one of the four L1s to access. <table border="1" data-bbox="267 1459 673 1669"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>GFX</td> </tr> <tr> <td>1h</td> <td>GPPSB</td> </tr> <tr> <td>2h</td> <td>GBIF</td> </tr> <tr> <td>3h</td> <td>INTGEN</td> </tr> <tr> <td>Fh-4h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	GFX	1h	GPPSB	2h	GBIF	3h	INTGEN	Fh-4h	Reserved
Bits	Description												
0h	GFX												
1h	GPPSB												
2h	GBIF												
3h	INTGEN												
Fh-4h	Reserved												
15:0	L1cfgIndex. Read-write. Reset: 0.												

D0F2xFC IOMMU L1 Config Data

See [D0F2xF8](#).

D0F2xFC_x00_L1sel[3:0] L1_PERF_CNTL

Bits	Description
31:24	L1PerfCountHi1 . Read-only. Reset: 0. read back of perf counter 1 bits 39:32
23:16	L1PerfCountHi0 . Read-only. Reset: 0. read back of perf counter 0 bits 39:32
15:8	L1PerfEvent1 . Read-write. Reset: 0. perf counter event 1
7:0	L1PerfEvent0 . Read-write. Reset: 0. perf counter event 0

D0F2xFC_x01_L1sel[3:0] L1_PERF_COUNT_0

Bits	Description
31:0	L1PerfCount0 . Read-only. Reset: 0. read back of perf counter 0 bits 31:0

D0F2xFC_x02_L1sel[3:0] L1_PERF_COUNT_1

Bits	Description
31:0	L1PerfCount1 . Read-only. Reset: 0. read back of perf counter 1 bits 31:0

D0F2xFC_x07_L1sel[3:0] L1_DEBUG_1

Bits	Description
31:18	Reserved.
17	L1NwEn . Read-write. Reset: 0. BIOS: 1. 1=Enable NW bit on ATS requests.
16:15	Reserved.
14	AtsPhysPageOverlapDis . Read-write. Reset: 0. BIOS: 1. 1=Prevent physical page overlap for ATS responses.
13	Reserved.
12	AtsSeqNumEn . Read-write. Reset: 0. BIOS: 1. 1=Enable logging of ATS sequence number.
11	SpecReqFilterEn . Read-write. Reset: 0. BIOS: 1. 1=Filter special requests in L1 work queue.
10:1	Reserved.
0	PhantomFuncEn . Read-write. Reset: 0. BIOS: See 2.12.2 . 1=Enable phantom function support.

D0F2xFC_x08_L1sel[3:0] L1_DEBUG_STATUS

Bits	Description
31:0	L1debugStatus . Read-only. Reset: 0. HW status bits reserved for ECOs

D0F2xFC_x0C_L1sel[3:0] L1_CNTRL_0

Bits	Description														
31	Reserved.														
30:28	L1VirtOrderQueues. Read-write. Reset: 0. BIOS: 4h. This field controls number of virtual queues in the L1 work queue. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>1</td> </tr> <tr> <td>1h</td> <td>2</td> </tr> <tr> <td>2h</td> <td>4</td> </tr> <tr> <td>3h</td> <td>8</td> </tr> <tr> <td>4h</td> <td>16</td> </tr> <tr> <td>7h-5h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	1	1h	2	2h	4	3h	8	4h	16	7h-5h	Reserved
Bits	Description														
0h	1														
1h	2														
2h	4														
3h	8														
4h	16														
7h-5h	Reserved														
27:24	L1Entries. Read-only; updated-by-hardware. Reset: 0. This field specifies the number of entries in each L1 cache as $2^{L1entries}$.														
23:22	Reserved.														
21:20	L1Banks. Read-only; updated-by-hardware. Reset: 1. This field specifies number of caches in L1.														
19:14	Reserved1.														
13:8	L2Credits. Read-write. Reset: 4h. This field controls credits for L1 to L2 interface.														
7:6	Reserved.														
5	ReplacementSel. Read-write. Reset: 0.														
4	Reserved0.														
3	CacheiwOnly. Read-write. Reset: 1. cache write only pages in L1														
2	CacheirOnly. Read-write. Reset: 1. cache read only pages in L1														
1	FragmentDis. Read-write. Reset: 0. disable variable page size support in L1 cache - only 4K pages														
0	UnfilterDis. Read-write. Reset: 0. disable unfiltering in L1 wq of aborted L2 requests														

D0F2xFC_x0D_L1sel[3:0] L1_CNTRL_1

Bits	Description
31	Reserved.
30	L1DebugCntrMode. Read-write. Reset: 0. Mode control for debug bus
29	Untrans2mFilteren. Read-write. Reset: 0. Enable filtering of requests on a 2M boundry instead of 4K
28	PretransNovaFilteren. Read-write. Reset: 0. When set, VA is not used for filtering pretrans requests
27	L1CacheSelInterleave. Read-write. Reset: 0. when set causes cache updates to toggle between multiple caches
26	L1CacheSelReqid. Read-write. Reset: 0. when set will allow the reqid to be used in hashing between multiple L1 caches

25:23	SelectTimeoutPulse. Read-write. Reset: 0.
22	L1CacheInvAllEn. Read-write. Reset: 0. Enables invalidation of entire cache when invalidation command is sent
21	L1orderEn. Read-write. Reset: 0. Enables strict ordering of all requests through L1
20	SndFilterDis. Read-write. Reset: 0. Disables filtering of requests to L2
19	AtsNobufferInsert. Read-write. Reset: 0. disables buffering of read completion data when inserting ats responses
18:14	WqEntrydis. Read-write. Reset: 0. Value indicates how many cache entries in L1 to disable
13	BlockL1Dis. Read-only. Reset: 0. N/A
12	L1DTEDis. Read-write. Reset: 0. Disables L1 caching of DTE
11	L1ParityEn. Read-write. Reset: 0.
10	L1CacheParityEn. Read-write. Reset: 0. Enables forced miss of L1 cache due to failed parity check
9	CacheByPass. Read-write. Reset: 0. Enables L1 cache bypass
8	VOQXorMode. Read-write. Reset: 0. N/A
7	Reserved.
6:4	VOQFuncBits. Read-write. Reset: 0. N/A
3	Reserved.
2:0	VOQPortBits. Read-write. Reset: 0. BIOS: See 2.12.2. When not 000b, enables virtual queue hashing using port id. This field controls the number of bits to use from portid for hashing.

D0F2xFC_x0E_L1sel[3:0] L1_CNTRL_2

Bits	Description
31:28	Reserved.
27:20	MsiHtRsvIntVector. Read-write. Reset: 00h. This field defines the interrupt vector used when an MSI interrupt is received that has a reserved DM field.
19:12	MsiHtRsvIntDestination. Read-write. Reset: FFh. This field defines the interrupt destination used when an MSI interrupt is received that has a reserved DM field.
11	Reserved.
10	MsiHtRsvIntDM. Read-write. Reset: 0. Defines the interrupt destination mode when an MSI interrupt is received that has a reserved DM field.
9	MsiHtRsvIntRqEio. Read-write. Reset: 0. Specifies the RQEOI state when an MSI interrupt is received that has a reserved DM field.
8:6	MsiHtRsvIntMt. Read-write. Reset: 011b. Specifies the message type used when an MSI interrupt is received that has a reserved DM field.
5:3	Reserved.
2	L1AbrtAtsDis. Read-write. Reset: 0. 1=Disable abort of ATS requests when IOMMU is disabled.
1	MsiToHtRemapDis. Read-write. Reset: 0. 1=Disable mapping of MSI to HT interrupts.
0	Reserved.

D0F2xFC_x0F_L1sel[3:0] L1_CNTRL_3

Bits	Description
31:0	AtsTlbinvPulseWidth . Read-write. Reset: C350h. sets the pulse width of the ats invalidation counters

D0F2xFC_x10_L1sel[3:0] L1_BANK_SEL_0

Bits	Description
31:16	Reserved.
15:0	L1cachebanksel0 . Read-write. Reset: 1. value is used to determine the virtual address bit that selects between the 2 banks of the L1 cache (if present). The bank is selected by bitwise ANDing this register against virtual address bits 19:12 and XORing the result

D0F2xFC_x11_L1sel[3:0] L1_BANK_DISABLE_0

Bits	Description
31:14	Reserved.
13:8	L1cachelinedis1 . Read-write. Reset: 0. sets the number of cache entries to disable in cache 1
7:6	Reserved.
5:0	L1cachelinedis0 . Read-write. Reset: 0. sets the number of cache entries to disable in cache 0

D0F2xFC_x20_L1sel[3:0] L1_WQ_STATUS_0**Table 121: Valid values for D0F2xFC_x20_L1sel[3:0]**

Value	Description
0h	Idle.
1h	Wait_L1.
2h	Wait_L2.
3h	Sending special request to L2.
4h	Waiting for completion of special request.
5h	Done.

Bits	Description
31:30	Reserved.
29:27	EntryStatus9 . See: EntryStatus0.
26:24	EntryStatus8 . See: EntryStatus0.
23:21	EntryStatus7 . See: EntryStatus0.
20:18	EntryStatus6 . See: EntryStatus0.
17:15	EntryStatus5 . See: EntryStatus0.

14:12	EntryStatus4 . See: EntryStatus0.
11:9	EntryStatus3 . See: EntryStatus0.
8:6	EntryStatus2 . See: EntryStatus0.
5:3	EntryStatus1 . See: EntryStatus0.
2:0	EntryStatus0 . Read-only. Reset: 0. See: Table 121 .

D0F2xFC_x21_L1sel[3:0] L1_WQ_STATUS_1

Bits	Description
31:30	Reserved.
29:27	EntryStatus19 . Read-only. Reset: 0.
26:24	EntryStatus18 . Read-only. Reset: 0.
23:21	EntryStatus17 . Read-only. Reset: 0.
20:18	EntryStatus16 . Read-only. Reset: 0.
17:15	EntryStatus15 . Read-only. Reset: 0.
14:12	EntryStatus14 . Read-only. Reset: 0.
11:9	EntryStatus13 . Read-only. Reset: 0.
8:6	EntryStatus12 . Read-only. Reset: 0.
5:3	EntryStatus11 . Read-only. Reset: 0.
2:0	EntryStatus10 . Read-only. Reset: 0.

D0F2xFC_x22_L1sel[3:0] L1_WQ_STATUS_2

Bits	Description
31:30	Reserved.
29:27	EntryStatus29 . Read-only. Reset: 0.
26:24	EntryStatus28 . Read-only. Reset: 0.
23:21	EntryStatus27 . Read-only. Reset: 0.
20:18	EntryStatus26 . Read-only. Reset: 0.
17:15	EntryStatus25 . Read-only. Reset: 0.
14:12	EntryStatus24 . Read-only. Reset: 0.
11:9	EntryStatus23 . Read-only. Reset: 0.
8:6	EntryStatus22 . Read-only. Reset: 0.
5:3	EntryStatus21 . Read-only. Reset: 0.
2:0	EntryStatus20 . Read-only. Reset: 0.

D0F2xFC_x23_L1sel[3:0] L1_WQ_STATUS_3

Bits	Description
31:16	Reserved.
15:8	InvalidationStatus . Read-only. Reset: 0. status of invalidation state machine
7:6	Reserved.
5:3	EntryStatus31 . Read-only. Reset: 0.
2:0	EntryStatus30 . Read-only. Reset: 0.

D0F2xFC_x32_L1sel[3:0] L1_CNTRL_4

Bits	Description
31:4	Reserved.
3	TlpprefixerrEn . Read-write. Reset: 0.
2	TimeoutPulseExtEn . Read-write. Reset: 0.
1	AtsMultipleL1toL2En . Read-write. Reset: 0. BIOS: 1.
0	AtsMultipleRespEn . Read-write. Reset: 0. BIOS: 1.

D0F2xFC_x33_L1sel[3:0] L1_CLKCNTRL_0

Bits	Description
31	L1L2ClkgateEn . Read-write. Reset: 0. BIOS: 1.
30:12	Reserved.
11	L1HostreqClkgateEn . Read-write. Reset: 0. BIOS: 0.
10	L1RegClkgateEn . Read-write. Reset: 0. BIOS: 1.
9	L1MemoryClkgateEn . Read-write. Reset: 0. BIOS: 1.
8	L1PerfClkgateEn . Read-write. Reset: 0. BIOS: 1.
7	L1DmaInputClkgateEn . Read-write. Reset: 0. BIOS: 1.
6	L1CpslvClkgateEn . Read-write. Reset: 0. BIOS: 1.
5	L1CacheClkgateEn . Read-write. Reset: 0. BIOS: 1.
4	L1DmaClkgateEn . Read-write. Reset: 0. BIOS: 1.
3:2	Reserved.
1:0	L1ClkgateLen . Read-write. Reset: 0.

D0F2xFC_x34_L1sel[3:0] L1_MEMPWRCNTRL_0

Bits	Description
31:24	L1MempwrTimer2 . Read-write. Reset: Fh.
23:16	L1MempwrTimer1 . Read-write. Reset: Fh.

15:8	L1MempwrTimer0. Read-write. Reset: Fh.
7:1	Reserved.
0	L1MempwrEn. Read-write. Reset: 0.

D0F2xFC_x35_L1sel[3:0] L1_MEMPWRCNTRL_1

Bits	Description
31:8	Reserved.
7:0	L1MempwrTimer3. Read-write. Reset: Fh.

D0F2xFC_x36_L1sel[3:0] L1_GUEST_ADDR_CNTRL

Bits	Description
31:8	L1GuestAddrMsk. Read-write. Reset: 0.
7:1	Reserved.
0	L1CanonicalErrEn. Read-write. Reset: 0.

D0F2xFC_x37_L1sel[3:0] L1_FEATURE_SUP_CNTRL

Bits	Description
31:2	Reserved.
1	L1PprSup. Read-write. Reset: 1.
0	L1EfrSup. Read-write. Reset: 1.

3.5 Device 1 Function 0 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space].

D1F0x00 Device/Vendor ID Register

Bits	Description
31:16	DeviceID: device ID. Value: Product-specific.
15:0	VendorID: vendor ID. Read-only. Reset: 1002h.

D1F0x04 Status/Command Register

Reset: 0010_0000h.

Bits	Description
31	ParityErrorDetected: detected parity error. Read; Write-1-to-clear. 1=Poisoned TLP received.
30	SignaledSystemError: signaled system error. Read; Write-1-to-clear. 1=A non-fatal or fatal error message was sent and SerrEn=1.
29	ReceivedMasterAbort: received master abort. Read; Write-1-to-clear. 1=A completion with an unsupported request completion status was received.
28	ReceivedTargetAbort: received target abort. Read; Write-1-to-clear. 1=A completion with completer abort completion status was received.
27	SignalTargetAbort: Signaled target abort. Read-only.
26:25	DevselTiming: DEVSEL# Timing. Read-only.
24	MasterDataPerr: master data parity error. Read; Write-1-to-clear. 1=ParityErrorEn=1 and either a poisoned completion was received or the device poisoned a write request.
23	FastBackCapable: fast back-to-back capable. Read-only.
22	UDFEn: UDF enable. Read-only.
21	PCI66En: 66 MHz capable. Read-only.
20	CapList: capability list. Read-only. 1=capability list supported.
19	IntStatus: interrupt status. Read-only. 1=INTx interrupt message pending.
18:11	Reserved.
10	IntDis: interrupt disable. Read-write. 1=INTx interrupt messages generation disabled.
9	FastB2BEn: fast back-to-back enable. Read-only.
8	SerrEn: System error enable. Read-write. 1=Enables reporting of non-fatal and fatal errors detected.
7	Stepping: Stepping control. Read-only.
6	ParityErrorEn: parity error response enable. Read-write.
5	PalSnoopEn: VGA palette snoop enable. Read-only.
4	MemWriteInvalidateEn: memory write and invalidate enable. Read-only.
3	SpecialCycleEn: special cycle enable. Read-only.

2	BusMasterEn: bus master enable. Read-write. 1=Memory and IO read and write request generation enabled.
1	MemAccessEn: IO access enable. Read-write. This bit controls if memory accesses targeting this device are accepted. 1=Enabled. 0=Disabled.
0	IoAccessEn: IO access enable. Read-write. This bit controls if IO accesses targeting this device are accepted. 1=Enabled. 0=Disabled.

D1F0x08 Class Code/Revision ID Register

Bits	Description
31:8	ClassCode. Value: 03_0000h.
7:0	RevID: revision ID. Value: Product-specific.

D1F0x0C Header Type Register

Reset: 0080_0000h.

Bits	Description
31:24	BIST. Read-only.
23:16	HeaderTypeReg. Read-only. The header type field indicates a header type 0 and that this is a multi-function device.
15:8	LatencyTimer. Read-only. These bits are fixed at their default value.
7:0	CacheLineSize. Read-write. This field specifies the system cache line size in units of double words.

D1F0x10 Graphic Memory Base Address

IF (D0F0x64_x1C[F064BarEn]==0) THEN Reset: 0000_0008h. ELSE Reset: 0000_000Ch. ENDIF.

Bits	Description
31:26	BaseAddr[31:26]: base address. Read-write. The amount of memory requested by the graphics memory BAR is controlled by D0F0x64_x1C[MemApSize].
25:4	BaseAddr[25:4]: base address. Read-only.
3	Pref: prefetchable. Read-only. 1=Prefetchable memory region.
2:1	Type: base address register type. Read-only. 00b=32-bit BAR. 10b=64-bit BAR.
0	MemSpace: memory space type. Read-only. 0=Memory mapped base address.

IF (D0F0x64_x1C[F064BarEn]==0) THEN

D1F0x14 Graphics IO Base Address

Reset: 0000_0001h.

Bits	Description
31:8	BaseAddr: base address. IF (D0F0x64_x1C[IoBarDis]==0) THEN Read-write. ELSE Read-only. ENDIF.

7:1	Reserved.
0	MemSpace: memory space type. Read-only. 1=IO mapped base address.

ELSE

D1F0x14 Graphics Memory Base Address 64

Reset: 0000_0000h.

Bits	Description
31:0	BaseAddr[63:32]: base address. Read-write.

ENDIF.

D1F0x18 Graphics Memory Mapped Registers Base Address

IF (D0F0x64_x1C[F064BarEn]==0) THEN Reset: 0000_0000h. ELSE Reset: 0000_0004h. ENDIF.

Bits	Description
31:18	BaseAddr[31:18]: base address. Read-write.
17:16	BaseAddr[17:16]: base address. IF (D0F0x64_x1C[RegApSize]==0) THEN Read-write. ELSE Read-only. ENDIF.
15:4	BaseAddr[15:4]: base address. Read-only.
3	Pref: prefetchable. Read-only. 0=Non-prefetchable memory region.
2:1	Type: base address register type. Read-only. 00b=32-bit BAR. 10b=64-bit BAR.
0	MemSpace: memory space type. Read-only. 0=Memory mapped base address.

IF (D0F0x64_x1C[F064BarEn]==0) THEN

D1F0x1C Base Address 3

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

ELSE

D1F0x1C Graphics Memory Mapped Registers Address 64

Reset: 0000_0000h.

Bits	Description
31:0	BaseAddr[63:32]: base address. Read-write.

ENDIF.

IF (D0F0x64_x1C[F064BarEn]==0) THEN

D1F0x20 Base Address 4

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

ELSE

D1F0x20 Graphics IO Base Address

Reset: 0000_0000h.

Bits	Description
31:8	BaseAddr: base address. IF (D0F0x64_x1C[IoBarDis]==0) THEN Read-write. ELSE Read-only. ENDIF.
7:1	Reserved.
0	MemSpace: memory space type. Read-only. 1=IO mapped base address.

ENDIF.

D1F0x24 Base Address 5

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F0x2C Subsystem and Subvendor ID RegisterReset: 0000_0000h. This register can be modified through [D1F0x4C](#)

Bits	Description
31:16	SubsystemID. Read-only.
15:0	SubsystemVendorID. Read-only.

D1F0x30 Expansion ROM Base Address

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F0x34 Capabilities Pointer

Reset: 0000_0050h.

Bits	Description
31:8	Reserved.
7:0	CapPtr: capabilities pointer. Read-only. Pointer to PM capability.

D1F0x3C Interrupt Line

Reset: 0000_01FFh.

Bits	Description
31:16	Reserved.
15:8	InterruptPin: interrupt pin. Read-only. This field identifies the legacy interrupt message the function uses.
7:0	InterruptLine: interrupt line. Read-write. This field contains the interrupt line routing information.

D1F0x4C Subsystem and Subvendor ID Mirror

Reset: 0000_0000h.

Bits	Description
31:16	SubsystemID. Read-write. This field sets the value in the corresponding field in D1F0x2C .
15:0	SubsystemVendorID. Read-write. This field sets the value in the corresponding field in D1F0x2C .

D1F0x50 Power Management Capability

Bits	Description
31:27	PmeSupport. Value: 0_0000b. Indicates that there is no PME support.
26	D2Support: D2 support. Value: 1. D2 is supported
25	D1Support: D1 support. Value: 1. D1 is supported
24:22	AuxCurrent: auxiliary current. Value: 0.
21	DevSpecificInit: device specific initialization. Value: 0. Indicates that there is no device specific initialization necessary.
20	Reserved.
19	PmeClock. Value: 0.
18:16	Version: version. Value: 011b.
15:8	NextPtr: next pointer. Value: 58h.
7:0	CapID: capability ID. Value: 01h. Indicates that the capability structure is a PCI power management data structure.

D1F0x54 Power Management Control and Status

Reset: 0000_0000h.

Bits	Description								
31:24	PmeData . Read-only.								
23	BusPwrEn . Read-only.								
22	B2B3Support . Read-only. B states are not supported.								
21:16	Reserved.								
15	PmeStatus : PME status . Read-only.								
14:13	DataScale : data scale . Read-only.								
12:9	DataSelect : data select . Read-only.								
8	PmeEn : PME# enable . Read-only.								
7:4	Reserved.								
3	NoSoftReset : no soft reset . Read-only. Software is required to re-initialize the function when returning from D3 _{hot} .								
2	Reserved.								
1:0	PowerState : power state . Read-write. This 2-bit field is used both to determine the current power state of the root port and to set the root port into a new power state. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>D0</td> </tr> <tr> <td>10b-01b</td> <td>Reserved.</td> </tr> <tr> <td>11b</td> <td>D3hot</td> </tr> </tbody> </table>	Bits	Definition	00b	D0	10b-01b	Reserved.	11b	D3hot
Bits	Definition								
00b	D0								
10b-01b	Reserved.								
11b	D3hot								

D1F0x58 PCI Express Capability

Bits	Description
31:30	Reserved.
29:25	IntMessageNum : interrupt message number . Value: 0. This field indicates which MSI vector is used for the interrupt message.
24	SlotImplemented : Slot implemented . Value: 0.
23:20	DeviceType : device type . Value: IF (D0F0x64_x1C[RcieEn]==1) THEN 9 ELSEIF (D0F0x64_x1C[F0NonlegacyDeviceTypeEn]==0) THEN 1 ELSE 0 ENDIF.
19:16	Version . Value: 2h.
15:8	NextPtr : next pointer . Value: IF (D0F0x64_x1C[MsiDis]==0) THEN A0h ELSE 00h ENDIF.
7:0	CapID : capability ID . Value: 10h.

D1F0x5C Device Capability

Bits	Description
31:29	Reserved.
28	FlrCapable : function level reset capability . Value: 0.

27:26	CapturedSlotPowerScale: captured slot power limit scale. Value: 0.
25:18	CapturedSlotPowerLimit: captured slot power limit value. Value: 0.
17:16	Reserved.
15	RoleBasedErrReporting: role-based error reporting. Value: 1.
14:12	Reserved.
11:9	L1AcceptableLatency: endpoint L1 Acceptable Latency. Value: 111b.
8:6	L0SAcceptableLatency: endpoint L0s Acceptable Latency. Value: 110b.
5	ExtendedTag: extended tag support. Value: 1. 8 bit tag support.
4:3	PhantomFunc: phantom function support. Value: 0. No phantom functions supported.
2:0	MaxPayloadSupport: maximum supported payload size. Value: 000b. 128 bytes max payload size.

D1F0x60 Device Control and Status

Reset: 0000_0810h.

Bits	Description
31:22	Reserved.
21	TransactionsPending: transactions pending. Read-only.
20	AuxPwr: auxiliary power. Read-only.
19	UsrDetected: unsupported request detected. Read; Write-1-to-clear. 1=Unsupported request received.
18	FatalErr: fatal error detected. Read; Write-1-to-clear. 1=Fatal error detected.
17	NonFatalErr: non-fatal error detected. Read; Write-1-to-clear. 1=Non-fatal error detected.
16	CorrErr: correctable error detected. Read; Write-1-to-clear. 1=Correctable error detected.
15	BridgeCfgRetryEn: bridge configuration retry enable. Read-only.
14:12	MaxRequestSize: maximum request size. Read-only.
11	NoSnoopEnable: enable no snoop. Read-write. 1=The device is permitted to set the No Snoop bit in requests.
10	AuxPowerPmEn: auxiliary power PM enable. Read-only. This capability is not implemented.
9	PhantomFuncEn: phantom functions enable. Read-only. Phantom functions are not supported.
8	ExtendedTagEn: extended tag enable. Read-write. 1=8-bit tag request tags. 0=5-bit request tag.
7:5	MaxPayloadSize: maximum supported payload size. Read-only. 000b=Indicates a 128 byte maximum payload size.
4	RelaxedOrdEn: relaxed ordering enable. Read-write. 1=The device is permitted to set the Relaxed Ordering bit.
3	UsrReportEn: unsupported request reporting enable. Read-write. 1=Enables signaling unsupported requests by sending error messages.
2	FatalErrEn: fatal error reporting enable. Read-write. 1=Enables sending ERR_FATAL message when a fatal error is detected.

1	NonFatalErrEn: non-fatal error reporting enable. Read-write. 1=Enables sending ERR_NONFATAL message when a non-fatal error is detected.
0	CorrErrEn: correctable error reporting enable. Read-write. 1=Enables sending ERR_CORR message when a correctable error is detected.

D1F0x64 Link Capability

Bits	Description
31:24	PortNumber: port number. Value: 0. This field indicates the PCI Express port number for the given PCI Express link.
23:22	Reserved.
21	LinkBWNotificationCap: link bandwidth notification capability. Value: 0b.
20	DLActiveReportingCapable: data link layer active reporting capability. Value: 0b.
19	SurpriseDownErrReporting: surprise down error reporting capability. Value: 0b.
18	ClockPowerManagement: clock power management. Value: 0b.
17:15	L1ExitLatency: L1 exit latency. Value: 010b.
14:12	L0sExitLatency: L0s exit latency. Value: 001b.
11:10	PMSupport: active state power management support. Value: 11b.
9:4	LinkWidth: maximum link width. Value: 0.
3:0	LinkSpeed: link speed. Value: 1.

D1F0x68 Link Control and Status

Reset: 1001_0000h.

Bits	Description
31	LinkAutonomousBWStatus: link autonomous bandwidth status. Read-only.
30	LinkBWManagementStatus: link bandwidth management status. Read-only.
29	DLActive: data link layer link active. Read-only. This bit indicates the status of the data link control and management state machine. Reads return a 1 to indicate the DL_Active state, otherwise 0 is returned.
28	SlotClockCfg: slot clock configuration. Read-only. 1=the root port uses the same clock that the platform provides.
27	LinkTraining: link training. Read-only. 1=Indicates that the physical layer link training state machine is in the configuration or recovery state, or that 1b was written to the RetrainLink bit but link training has not yet begun. Hardware clears this bit when the link training state machine exits the configuration/recovery state.
26	Reserved.
25:20	NegotiatedLinkWidth: negotiated link width. Read-only. This field indicates the negotiated width of the given PCI Express link.
19:16	LinkSpeed: link speed. Read-only.
15:12	Reserved.
11	LinkAutonomousBWIntEn: link autonomous bandwidth interrupt enable. Read-only.

10	LinkBWManagementEn: link bandwidth management interrupt enable. Read-only.												
9	HWAutonomousWidthDisable: hardware autonomous width disable. Read-write. 1=Hardware not allowed to change the link width except to correct unreliable link operation by reducing link width.												
8	ClockPowerManagementEn: clock power management enable. Read-write.												
7	ExtendedSync: extended sync. Read-write. 1=Forces the transmission of additional ordered sets when exiting the L0s state and when in the recovery state.												
6	CommonClockCfg: common clock configuration. Read-write. 1=Indicates that the root port and the component at the opposite end of this Link are operating with a distributed common reference clock. 0=Indicates that the upstream port and the component at the opposite end of this Link are operating with asynchronous reference clock.												
5	RetrainLink: retrain link. Read-only. This bit does not apply to endpoints.												
4	LinkDis: link disable. Read-only. This bit does not apply to endpoints.												
3	ReadCplBoundary: read completion boundary. Read-only. 0=64 byte read completion boundary.												
2	Reserved.												
1:0	PmControl: active state power management enable. Read-write. This field controls the level of ASPM supported on the given PCI Express link. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Disabled.</td> <td>10b</td> <td>L1 Entry Enabled.</td> </tr> <tr> <td>01b</td> <td>L0s Entry Enabled.</td> <td>11b</td> <td>L0s and L1 Entry Enabled.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	00b	Disabled.	10b	L1 Entry Enabled.	01b	L0s Entry Enabled.	11b	L0s and L1 Entry Enabled.
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>										
00b	Disabled.	10b	L1 Entry Enabled.										
01b	L0s Entry Enabled.	11b	L0s and L1 Entry Enabled.										

D1F0x7C Device Capability 2

Reset: 0000_0000h.

Bits	Description
31:5	Reserved.
4	CplTimeoutDisSup: completion timeout disable supported. Read-only.
3:0	CplTimeoutRangeSup: completion timeout range supported. Read-only.

D1F0x80 Device Control and Status 2

Reset: 0000_0000h.

Bits	Description
31:5	Reserved.
4	CplTimeoutDis: completion timeout disable. Read-only.
3:0	CplTimeoutValue: completion timeout range supported. Read-only.

D1F0x88 Link Control and Status 2

Reset: 0000_0001h.

Bits	Description
31:17	Reserved.

16	CurDeemphasisLevel: current deemphasis level. Read-only. 1=-3.5 dB. 0=-6 dB.
15:13	Reserved.
12	ComplianceDeemphasis: compliance deemphasis. Read-write. This bit defines the deemphasis level used in compliance mode. 1=-3.5 dB. 0=-6 dB.
11	ComplianceSOS: compliance SOS. Read-write. 1=The device transmits skip ordered sets in between the modified compliance pattern.
10	EnterModCompliance: enter modified compliance. Read-write. 1=The device transmits modified compliance pattern.
9:7	XmitMargin: transmit margin. Read-write. This field controls the non-deemphasized voltage level at the transmitter pins.
6	SelectableDeemphasis: selectable deemphasis. Read-only.
5	HwAutonomousSpeedDisable: hardware autonomous speed disable. Read-write. 1=Disables hardware generated link speed changes.
4	EnterCompliance: enter compliance. Read-write. 1=Force link to enter compliance mode.
3:0	TargetLinkSpeed: target link speed. Read-write. This field defines the upper limit of the link operational speed.

D1F0xA0 MSI Capability

Bits	Description
31:24	Reserved.
23	Msi64bit: MSI 64 bit capability. Read-only. Value: D0F0x64_x46[Msi64bitEn]. 1=The device is capable of sending 64-bit MSI messages. 0=The device is not capable of sending a 64-bit message address.
22:20	MsiMultiEn: MSI multiple message enable. Read-write. Reset: 000b. Software writes to this field to indicate the number of allocated vectors (equal to or less than the number of requested vectors). When MSI is enabled, a function is allocated at least 1 vector.
19:17	MsiMultiCap: MSI multiple message capability. Read-only. Reset: 000b. 000b=The device is requesting one vector.
16	MsiEn: MSI enable. Read-write. Reset: 0. 1=MSI generation is enabled and INTx generation is disabled. 0=MSI generation disabled and INTx generation is enabled.
15:8	NextPtr: next pointer. Read-only. Reset:00h
7:0	CapID: capability ID. Read-only. Reset: 05h. 05h=MSI capability structure.

D1F0xA4 MSI Message Address Low

Reset: 0000_0000h.

Bits	Description
31:2	MsiMsgAddrLo: MSI message address. Read-write. This register specifies the dword aligned address for the MSI memory write transaction.
1:0	Reserved.

IF (DOF0x64_x46[Msi64bitEn]==0) THEN

D1F0xA8 MSI Message Data

Reset: 0000_0000h.

Bits	Description
31:16	Reserved.
15:0	MsiData: MSI message data. Read-write. This register specifies lower 16 bits of data for the MSI memory write transaction. The upper 16 bits are always 0.

ELSE

D1F0xA8 MSI Message Address High

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7:0	MsiMsgAddrHi: MSI message address. Read-write. This register specifies the upper 8-bits of the MSI address in 64 bit MSI mode.

D1F0xAC MSI Message Data

Reset: 0000_0000h.

Bits	Description
31:16	Reserved.
15:0	MsiData: MSI message data. Read-write. This register specifies lower 16 bits of data for the MSI memory write transaction. The upper 16 bits are always 0.

ENDIF.

D1F0x100 Vendor Specific Enhanced Capability

Reset: 0061_000Bh.

Bits	Description
31:20	NextPtr: next pointer. Read-only.
19:16	CapVer: capability version. Read-only.
15:0	CapID: capability ID. Read-only.

D1F0x104 Vendor Specific Header

Reset: 0101_0001h.

Bits	Description
31:20	VsecLen: vendor specific enhanced next pointer. Read-only.

19:16	VsecRev: vendor specific enhanced capability version. Read-only.
15:0	VsecID: vendor specific enhanced capability ID. Read-only.

D1F0x108 Vendor Specific 1

Reset: 0000_0000h.

Bits	Description
31:0	Scratch: scratch. Read-write.

D1F0x10C Vendor Specific 2

Reset: 0000_0000h.

Bits	Description
31:0	Scratch: scratch. Read-write.

3.6 Device 1 Function 1 (Audio Controller) Configuration Registers

See 3.1 [Register Descriptions and Mnemonics] for a description of the register naming convention. See 2.7 [Configuration Space] for details about how to access this space.

D1F1x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Reset: 9902h.
15:0	VendorID: vendor ID. Read-only. Reset: 1002h.

D1F1x04 Status/Command

Reset: 0010_0000h.

Bits	Description
31	ParityErrorDetected: detected parity error. Read; Write-1-to-clear. 1=Poisoned TLP received.
30	SignaledSystemError: signaled system error. Read; Write-1-to-clear. 1=A non-fatal or fatal error message was sent and SerrEn=1.
29	ReceivedMasterAbort: received master abort. Read; Write-1-to-clear. 1=A completion with an unsupported request completion status was received.
28	ReceivedTargetAbort: received target abort. Read; Write-1-to-clear. 1=A completion with completer abort completion status was received.
27	SignalTargetAbort: Signaled target abort. Read-only.
26:25	DevselTiming: DEVSEL# Timing. Read-only.
24	MasterDataPerr: master data parity error. Read; Write-1-to-clear. 1=ParityErrorEn=1 and either a poisoned completion was received or the device poisoned a write request.
23	FastBackCapable: fast back-to-back capable. Read-only.
22	UDFEn: UDF enable. Read-only.
21	PCI66En: 66 MHz capable. Read-only.
20	CapList: capability list. Read-only. 1=capability list supported.
19	IntStatus: interrupt status. Read-only. 1=INTx interrupt message pending.
18:11	Reserved.
10	IntDis: interrupt disable. Read-write. 1=INTx interrupt messages generation disabled.
9	FastB2BEn: fast back-to-back enable. Read-only.
8	SerrEn: System error enable. Read-write. 1=Enables reporting of non-fatal and fatal errors detected.
7	Stepping: Stepping control. Read-only.
6	ParityErrorEn: parity error response enable. Read-write.
5	PalSnoopEn: VGA palette snoop enable. Read-only.
4	MemWriteInvalidateEn: memory write and invalidate enable. Read-only.
3	SpecialCycleEn: special cycle enable. Read-only.

2	BusMasterEn: bus master enable. Read-write. 1=Memory and IO read and write request generation enabled.
1	MemAccessEn: IO access enable. Read-write. This bit controls if memory accesses targeting this device are accepted. 1=Enabled. 0=Disabled.
0	IoAccessEn: IO access enable. Read-write. This bit controls if IO accesses targeting this device are accepted. 1=Enabled. 0=Disabled.

D1F1x08 Class Code/Revision ID

Reset: 0403_0000h.

Bits	Description
31:8	ClassCode. Read-only.
7:0	RevID: revision ID. Read-only.

D1F1x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:24	BIST. Read-only. These bits are fixed at their default values.
23:16	HeaderTypeReg. Read-only. 80h=Type 0 multi-function device.
15:8	LatencyTimer. Read-only. These bits are fixed at their default value.
7:0	CacheLineSize. Read-write. This field specifies the system cache line size in units of double words.

D1F1x10 Audio Registers Base Address

Reset: 0000_0000h.

Bits	Description
31:14	BaseAddr: base address. Read-write.
13:4	Reserved.
3	Pref: prefetchable. Read-only. 0=Non-prefetchable memory region.
2:1	Type: base address register type. Read-only. 00b=32-bit base address register.
0	MemSpace: memory space type. Read-only. 0=Memory mapped base address.

D1F1x14 Base Address 1

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F1x18 Base Address 2

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F1x1C Base Address 3

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F1x20 Base Address 4

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F1x24 Base Address 5

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F1x2C Subsystem and Subvendor IDReset: 0000_0000h. This register can be modified through [D1F1x4C](#).

Bits	Description
31:16	SubsystemID . Read-only.
15:0	SubsystemVendorID . Read-only.

D1F1x30 Expansion ROM Base Address

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D1F1x34 Capabilities Pointer

Reset: 0000_0050h.

Bits	Description
31:8	Reserved.
7:0	CapPtr: capabilities pointer. Read-only. Pointer to PM capability.

D1F1x3C Interrupt Line

Reset: 0000_02FFh.

Bits	Description
31:16	Reserved.
15:8	InterruptPin: interrupt pin. Read-only. This field identifies the legacy interrupt message the function uses.
7:0	InterruptLine: interrupt line. Read-write. This field contains the interrupt line routing information.

D1F1x4C Subsystem and Subvendor ID Mirror

Reset: 0000_0000h.

Bits	Description
31:16	SubsystemID. Read-write. This field sets the value in the corresponding field in D1F1x2C .
15:0	SubsystemVendorID. Read-write. This field sets the value in the corresponding field in D1F1x2C .

D1F1x50 Power Management Capability

Bits	Description
31:27	PmeSupport. Value: 0_0000b. Indicates that there is no PME support.
26	D2Support: D2 support. Value: 1. D2 is supported
25	D1Support: D1 support. Value: 1. D1 is supported
24:22	AuxCurrent: auxiliary current. Value: 0.
21	DevSpecificInit: device specific initialization. Value: 0. Indicates that there is no device specific initialization necessary.
20	Reserved.
19	PmeClock. Value: 0.
18:16	Version: version. Value: 011b.
15:8	NextPtr: next pointer. Value: 58h.
7:0	CapID: capability ID. Value: 01h. Indicates that the capability structure is a PCI power management data structure.

D1F1x54 Power Management Control and Status

Reset: 0000_0000h.

Bits	Description								
31:24	PmeData . Read-only.								
23	BusPwrEn . Read-only.								
22	B2B3Support . Read-only. B states are not supported.								
21:16	Reserved.								
15	PmeStatus : PME status . Read-only.								
14:13	DataScale : data scale . Read-only.								
12:9	DataSelect : data select . Read-only.								
8	PmeEn : PME# enable . Read-only.								
7:4	Reserved.								
3	NoSoftReset : no soft reset . Read-only. Software is required to re-initialize the function when returning from D3 _{hot} .								
2	Reserved.								
1:0	PowerState : power state . Read-write. This 2-bit field is used both to determine the current power state of the root port and to set the root port into a new power state. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>D0</td> </tr> <tr> <td>10b-01b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>D3hot</td> </tr> </tbody> </table>	Bits	Definition	00b	D0	10b-01b	Reserved	11b	D3hot
Bits	Definition								
00b	D0								
10b-01b	Reserved								
11b	D3hot								

D1F1x58 PCI Express Capability

Bits	Description
31:30	Reserved.
29:25	IntMessageNum : interrupt message number . Value: 0. This field indicates which MSI vector is used for the interrupt message.
24	SlotImplemented : Slot implemented . Value: 0.
23:20	DeviceType : device type . Value: IF (D0F0x64_x1C[AudioNonlegacyDeviceTypeEn]==0) THEN 1. ELSE 0. ENDIF.
19:16	Version . Value: 2h.
15:8	NextPtr : next pointer . Value: IF (D0F0x64_x1C[MsiDis]==0) THEN A0h. ELSE 00h. ENDIF.
7:0	CapID : capability ID . Value: 10h.

D1F1x5C Device Capability

Bits	Description
31:29	Reserved.
28	FlrCapable : function level reset capability . Value: 0.

27:26	CapturedSlotPowerScale: captured slot power limit scale. Value: 0.
25:18	CapturedSlotPowerLimit: captured slot power limit value. Value: 0.
17:16	Reserved.
15	RoleBasedErrReporting: role-based error reporting. Value: 1.
14:12	Reserved.
11:9	L1AcceptableLatency: endpoint L1 Acceptable Latency. Value: 111b.
8:6	L0SAcceptableLatency: endpoint L0s Acceptable Latency. Value: 110b.
5	ExtendedTag: extended tag support. Value: 1. 8 bit tag support.
4:3	PhantomFunc: phantom function support. Value: 0. No phantom functions supported.
2:0	MaxPayloadSupport: maximum supported payload size. Value: 000b. 128 bytes max payload size.

D1F1x60 Device Control and Status

Reset: 0000_0810h.

Bits	Description
31:22	Reserved.
21	TransactionsPending: transactions pending. Read-only.
20	AuxPwr: auxiliary power. Read-only.
19	UsrDetected: unsupported request detected. Read; Write-1-to-clear. 1=Unsupported request received.
18	FatalErr: fatal error detected. Read; Write-1-to-clear. 1=Fatal error detected.
17	NonFatalErr: non-fatal error detected. Read; Write-1-to-clear. 1=Non-fatal error detected.
16	CorrErr: correctable error detected. Read; Write-1-to-clear. 1=Correctable error detected.
15	BridgeCfgRetryEn: bridge configuration retry enable. Read-only.
14:12	MaxRequestSize: maximum request size. Read-only. 0=The root port never generates read requests with size exceeding 128 bytes.
11	NoSnoopEnable: enable no snoop. Read-write. 1=the device is permitted to set the No Snoop bit in requests.
10	AuxPowerPmEn: auxiliary power PM enable. Read-only. This capability is not implemented.
9	PhantomFuncEn: phantom functions enable. Read-only. Phantom functions are not supported.
8	ExtendedTagEn: extended tag enable. Read-write. 1=8-bit tag request tags. 0=5-bit request tag.
7:5	MaxPayloadSize: maximum supported payload size. Read-only. 000b=Indicates a 128 byte maximum payload size.
4	RelaxedOrdEn: relaxed ordering enable. Read-write. 1=The device is permitted to set the Relaxed Ordering bit.
3	UsrReportEn: unsupported request reporting enable. Read-write. 1=Enables signaling unsupported requests by sending error messages.
2	FatalErrEn: fatal error reporting enable. Read-write. 1=Enables sending ERR_FATAL message when a fatal error is detected.

1	NonFatalErrEn: non-fatal error reporting enable. Read-write. 1=Enables sending ERR_NONFATAL message when a non-fatal error is detected.
0	CorrErrEn: correctable error reporting enable. Read-write. 1=Enables sending ERR_CORR message when a correctable error is detected.

D1F1x64 Link Capability

Bits	Description
31:24	PortNumber: port number. Value: 0. This field indicates the PCI Express port number for the given PCI Express link.
23:22	Reserved.
21	LinkBWNotificationCap: link bandwidth notification capability. Value: 0b.
20	DIActiveReportingCapable: data link layer active reporting capability. Value: 0b.
19	SurpriseDownErrReporting: surprise down error reporting capability. Value: 0b.
18	ClockPowerManagement: clock power management. Value: 0b.
17:15	L1ExitLatency: L1 exit latency. Value: 010b.
14:12	L0sExitLatency: L0s exit latency. Value: 001b.
11:10	PMSupport: active state power management support. Value: 11b.
9:4	LinkWidth: maximum link width. Value: 0.
3:0	LinkSpeed: link speed. Value: 1.

D1F1x68 Link Control and Status

Reset: 1001_0000h.

Bits	Description
31	LinkAutonomousBWStatus: link autonomous bandwidth status. Read-only.
30	LinkBWManagementStatus: link bandwidth management status. Read-only.
29	DIActive: data link layer link active. Read-only. This bit indicates the status of the data link control and management state machine. Reads return a 1 to indicate the DL_Active state, otherwise 0 is returned.
28	SlotClockCfg: slot clock configuration. Read-only. 1=the root port uses the same clock that the platform provides.
27	LinkTraining: link training. Read-only. 1=Indicates that the physical layer link training state machine is in the configuration or recovery state, or that 1b was written to the RetrainLink bit but link training has not yet begun. Hardware clears this bit when the link training state machine exits the configuration/recovery state.
26	Reserved.
25:20	NegotiatedLinkWidth: negotiated link width. Read-only. This field indicates the negotiated width of the given PCI Express link.
19:16	LinkSpeed: link speed. Read-only. 0001b: 2.5 Gb/s. 0010b: 5 Gb/s

15:12	Reserved.												
11	LinkAutonomousBWIntEn: link autonomous bandwidth interrupt enable. Read-only.												
10	LinkBWManagementEn: link bandwidth management interrupt enable. Read-only.												
9	HWAutonomousWidthDisable: hardware autonomous width disable. Read-write. 1=Hardware not allowed to change the link width except to correct unreliable link operation by reducing link width.												
8	ClockPowerManagementEn: clock power management enable. Read-write.												
7	ExtendedSync: extended sync. Read-write. 1=Forces the transmission of additional ordered sets when exiting the L0s state and when in the recovery state.												
6	CommonClockCfg: common clock configuration. Read-write. 1=Indicates that the root port and the component at the opposite end of this Link are operating with a distributed common reference clock. 0=Indicates that the upstream port and the component at the opposite end of this Link are operating with asynchronous reference clock.												
5	RetrainLink: retrain link. Read-only. This bit does not apply to endpoints.												
4	LinkDis: link disable. Read-only. This bit does not apply to endpoints.												
3	ReadCplBoundary: read completion boundary. Read-only. 0=64 byte read completion boundary.												
2	Reserved.												
1:0	PmControl: active state power management enable. Read-write. This field controls the level of ASPM supported on the given PCI Express link. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Disabled.</td> <td>10b</td> <td>L1 Entry Enabled.</td> </tr> <tr> <td>01b</td> <td>L0s Entry Enabled.</td> <td>11b</td> <td>L0s and L1 Entry Enabled.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	00b	Disabled.	10b	L1 Entry Enabled.	01b	L0s Entry Enabled.	11b	L0s and L1 Entry Enabled.
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>										
00b	Disabled.	10b	L1 Entry Enabled.										
01b	L0s Entry Enabled.	11b	L0s and L1 Entry Enabled.										

D1F1x7C Device Capability 2

Reset: 0000_0000h.

Bits	Description
31:5	Reserved.
4	CplTimeoutDisSup: completion timeout disable supported. Read-only.
3:0	CplTimeoutRangeSup: completion timeout range supported. Read-only.

D1F1x80 Device Control and Status 2

Reset: 0000_0000h.

Bits	Description
31:5	Reserved.
4	CplTimeoutDis: completion timeout disable. Read-only.
3:0	CplTimeoutValue: completion timeout range supported. Read-only.

D1F1x84 Link Capability 2

Bits	Description
31:0	Reserved.

D1F1x88 Link Control and Status 2

Reset: 0000_0001h.

Bits	Description						
31:17	Reserved.						
16	CurDeemphasisLevel: current deemphasis level. Read-only. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>-3.5 dB</td> </tr> <tr> <td>0</td> <td>-6 dB</td> </tr> </tbody> </table>	Bit	Description	1	-3.5 dB	0	-6 dB
Bit	Description						
1	-3.5 dB						
0	-6 dB						
15:13	Reserved.						
12	ComplianceDeemphasis: compliance deemphasis. Read-write. This bit defines the deemphasis level used in compliance mode. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>-3.5 dB</td> </tr> <tr> <td>0</td> <td>-6 dB</td> </tr> </tbody> </table>	Bit	Description	1	-3.5 dB	0	-6 dB
Bit	Description						
1	-3.5 dB						
0	-6 dB						
11	ComplianceSOS: compliance SOS. Read-write. 1=The device transmits skip ordered sets in between the modified compliance pattern.						
10	EnterModCompliance: enter modified compliance. Read-write. 1=The device transmits modified compliance pattern.						
9:7	XmitMargin: transmit margin. Read-write. This field controls the non-deemphasized voltage level at the transmitter pins.						
6	SelectableDeemphasis: selectable deemphasis. Read-only.						
5	HwAutonomousSpeedDisable: hardware autonomous speed disable. Read-write. 1=Disables hardware generated link speed changes.						
4	EnterCompliance: enter compliance. Read-write. 1=Force link to enter compliance mode.						
3:0	TargetLinkSpeed: target link speed. Read-write. This fields defines the upper limit of the link operational speed.						

D1F1xA0 MSI Capability

Bits	Description
31:24	Reserved.
23	Msi64bit: MSI 64 bit capability. Read-only. Value: D0F0x64_x46[Msi64bitEn]. 1=The device is capable of sending 64-bit MSI messages. 0=The device is not capable of sending a 64-bit message address.
22:20	MsiMultiEn: MSI multiple message enable. Read-write. Reset: 000b. Software writes to this field to indicate the number of allocated vectors (equal to or less than the number of requested vectors). When MSI is enabled, a function is allocated at least 1 vector.

19:17	MsiMultiCap: MSI multiple message capability. Read-only. Reset:000b. 000b=The device is requesting one vector.
16	MsiEn: MSI enable. Read-write. Reset: 0. 1=MSI generation is enabled and INTx generation is disabled. 0=MSI generation disabled and INTx generation is enabled.
15:8	NextPtr: next pointer. Read-only. Reset: 00h
7:0	CapID: capability ID. Read-only. Reset: 05h. 05h=MSI capability structure.

D1F1xA4 MSI Message Address Low

Reset: 0000_0000h.

Bits	Description
31:2	MsiMsgAddrLo: MSI message address. Read-write. This register specifies the dword aligned address for the MSI memory write transaction.
1:0	Reserved.

IF (DOF0x64_x46[Msi64bitEn]==0) THEN

D1F1xA8 MSI Message Data

Reset: 0000_0000h.

Bits	Description
31:16	Reserved.
15:0	MsiData: MSI message data. Read-write. This register specifies lower 16 bits of data for the MSI memory write transaction. The upper 16 bits are always 0.

ELSE

D1F1xA8 MSI Message Address High

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7:0	MsiMsgAddrHi: MSI message address. Read-write. This register specifies the upper 8-bits of the MSI address in 64 bit MSI mode.

D1F1xAC MSI Message Data

Reset: 0000_0000h.

Bits	Description
31:16	Reserved.
15:0	MsiData: MSI message data. Read-write. This register specifies lower 16 bits of data for the MSI memory write transaction. The upper 16 bits are always 0.

ENDIF.

D1F1x100 Vendor Specific Enhanced Capability

Reset: 0111_000Bh.

Bits	Description
31:20	NextPtr: next pointer. Read-only.
19:16	CapVer: capability version. Read-only.
15:0	CapID: capability ID. Read-only.

D1F1x104 Vendor Specific Header

Reset: 0101_0001h.

Bits	Description
31:20	VsecLen: vendor specific enhanced next pointer. Read-only.
19:16	VsecRev: vendor specific enhanced capability version. Read-only.
15:0	VsecID: vendor specific enhanced capability ID. Read-only.

D1F1x108 Vendor Specific 1

Reset: 0000_0000h.

Bits	Description
31:0	Scratch: scratch. Read-write.

D1F1x10C Vendor Specific 2

Reset: 0000_0000h.

Bits	Description
31:0	Scratch: scratch. Read-write.

3.7 Device [8:2] Function 0 (Root Port) Configuration Registers

See 3.1 [Register Descriptions and Mnemonics] for a description of the register naming convention. See 2.7 [Configuration Space] for details about how to access this space.

D[8:2]F0x00 Device/Vendor ID Register

Table 122: Reset mapping for D[8:2]F0x00.

Register	Reset	Register	Reset
D2F0x00	1412_1022h.	D6F0x00	1416_1022h.
D3F0x00	1413_1022h.	D7F0x00	1417_1022h.
D4F0x00	1414_1022h.	D8F0x00	1418_1022h.
D5F0x00	1415_1022h.	-	-

Bits	Description
31:16	DeviceID: device ID. Read-only.
15:0	VendorID: vendor ID. Read-only.

D[8:2]F0x04 Status/Command Register

Reset: 0010_0000h.

Bits	Description
31	ParityErrorDetected: detected parity error. Read; Write-1-to-clear; updated-by-hardware.
30	SignaledSystemError: signaled system error. Read; Write-1-to-clear; updated-by-hardware. 1=System error signalled.
29	ReceivedMasterAbort: received master abort. Read; Write-1-to-clear; updated-by-hardware.
28	ReceivedTargetAbort: received target abort. Read; Write-1-to-clear; updated-by-hardware.
27	SignalTargetAbort: signaled target abort. Read; Write-1-to-clear; updated-by-hardware.
26:25	DevselTiming: DEVSEL# Timing. Read-only.
24	DataPerr: data parity error. Read; Write-1-to-clear; updated-by-hardware.
23	FastBackCapable: fast back-to-back capable. Read-only.
22	UDFEn: UDF enable. Read-only.
21	PCI66En: 66 MHz capable. Read-only.
20	CapList: capability list. Read-only. 1= Capability list present.
19	IntStatus: interrupt status. Read-only. 1=An INTx interrupt Message is pending in the device.
18:11	Reserved.
10	IntDis: interrupt disable. Read-write.
9	FastB2BEn: fast back-to-back enable. Read-only.
8	SerrEn: system error enable. Read-write. 1=System error reporting enabled.
7	Stepping: Stepping control. Read-only.
6	ParityErrorEn: parity error response enable. Read-write.
5	PalSnoopEn: VGA palette snoop enable. Read-only.

4	MemWriteInvalidateEn: memory write and invalidate enable. Read-only.
3	SpecialCycleEn: special cycle enable. Read-only.
2	BusMasterEn: bus master enable. Read-write.
1	MemAccessEn: memory access enable. Read-write. This bit controls if memory accesses targeting this device are accepted or not. 1=Enabled. 0=Disabled.
0	IoAccessEn: IO access enable. Read-write. This bit controls if IO accesses targeting this device are accepted or not. 1=Enabled. 0=Disabled.

D[8:2]F0x08 Class Code/Revision ID Register

Reset: 0604_00xxh.

Bits	Description
31:8	ClassCode. Read-only. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only.

D[8:2]F0x0C Header Type Register

Reset: 0001_0000h.

Bits	Description
31:24	BIST. Read-only. These bits are fixed at their default values.
23:16	HeaderTypeReg. Read-only. These bits are fixed at their default values. The header type field indicates a header type 1 and that there is only one function present in this device.
15:8	LatencyTimer. Read-only. This field does not control any hardware.
7:0	CacheLineSize. Read-write.

D[8:2]F0x18 Bus Number and Secondary Latency Register

Reset: 0000_0000h.

Bits	Description
31:24	SecondaryLatencyTimer: secondary latency timer. Read-only. This field is always 0.
23:16	SubBusNumber: subordinate number. Read-write. This field contains the highest-numbered bus that exists on the secondary side of the bridge.
15:8	SecondaryBus: secondary bus number. Read-write. This field defines the bus number of the secondary bus interface.
7:0	PrimaryBus: primary bus number. Read-write. This field defines the bus number of the primary bus interface.

D[8:2]F0x1C IO Base and Secondary Status Register

Reset: 0000_0101h.

Bits	Description
31	ParityErrorDetected: detected parity error. Read; Write-1-to-clear. A Poisoned TLP was received regardless of the state of the D[8:2]F0x04[ParityErrorEn].
30	ReceivedSystemError: signaled system error. Read; Write-1-to-clear. 1=A System Error was detected.
29	ReceivedMasterAbort: received master abort. Read; Write-1-to-clear. 1=A CPU transaction is terminated due to a master-abort.
28	ReceivedTargetAbort: received target abort. Read; Write-1-to-clear. 1=A CPU transaction (except for a special cycle) is terminated due to a target-abort.
27	SignalTargetAbort: signaled target abort. Read; Write-1-to-clear.
26:25	DevselTiming: DEVSEL# Timing. Read-only.
24	MasterDataPerr: master data parity error. Read; Write-1-to-clear. 1=The link received a poisoned or poisoned a downstream write and D[8:2]F0x3C[ParityResponseEn]=1.
23	FastBackCapable: fast back-to-back capable. Read-only.
22	UDFEn: UDF enable. Read-only.
21	PCI66En: 66 MHz capable. Read-only.
20	CapList: capability list. Read-only.
19:16	Reserved.
15:12	IOLimit[15:12]. Read-write. Lower part of the limit address. Upper part is defined in D[8:2]F0x30.
11:8	Reserved.
7:4	IOBase[15:12]. Read-write. Lower part of the base address. Upper part is defined in D[8:2]F0x30.
3:0	Reserved.

D[8:2]F0x20 Memory Limit and Base Register

Reset: 0000_0000h.

Bits	Description
31:20	MemLimit. Read-write.
19:16	Reserved.
15:4	MemBase. Read-write.
3:0	Reserved.

D[8:2]F0x24 Prefetchable Memory Limit and Base Register

Reset: 0001_0001h.

Bits	Description
31:20	PrefMemLimit. Read-write. Lower part of the limit address. Upper part is defined in D[8:2]F0x2C.
19:16	PrefMemLimitR. Read-only. 1=64 bit memory address decoder.

15:4	PrefMemBase[31:20] . Read-write. Lower part of the base address. Upper part is defined in D[8:2]F0x28 .
3:0	PrefMemBaseR . Read-only. 1=64 bit memory address decoder.

D[8:2]F0x28 Prefetchable Memory Base High Register

Reset: 0000_0000h.

Bits	Description
31:0	PrefMemBase[63:32] . Read-write. Upper part of the base address. Lower part is defined in D[8:2]F0x24 .

D[8:2]F0x2C Prefetchable Memory Limit High Register

Reset: 0000_0000h.

Bits	Description
31:0	PrefMemLimit[63:32] . Read-write. Upper part of the limit address. Lower part is defined in D[8:2]F0x24 .

D[8:2]F0x30 IO Base and Limit High Register

Reset: 0000_0000h.

Bits	Description
31:16	IOLimit[31:16] . Read-write. Upper part of the limit address. Lower part is defined in D[8:2]F0x1C .
15:0	IOBase[31:16] . Read-write. Upper part of the base address. Lower part is defined in D[8:2]F0x1C .

D[8:2]F0x34 Capabilities Pointer Register

Reset: 0000_0050h.

Bits	Description
31:8	Reserved.
7:0	CapPtr: capabilities pointer . Read-only. Pointer to PM capability.

D[8:2]F0x3C Bridge Control Register

Reset: 0000_00FFh.

Bits	Description
31:24	Reserved.
23	FastB2BCap: Fast back-to-back capability . Read-only.
22	SecondaryBusReset: Secondary bus reset . Read-write. Setting this bit triggers a hot reset on the corresponding PCI Express Port.
21	MasterAbortMode: Master abort mode . Read-only.

20	Vga16En: VGA IO 16 bit decoding enable. Read-write. 1= Address bits 15:10 for VGA IO cycles are decoded. 0=Address bits 15:10 for VGA IO cycles are ignored.
19	VgaEn: VGA enable. Read-write. Affects the response by the bridge to compatible VGA addresses. When it is set, the bridge decodes and forwards the following accesses on the primary interface to the secondary interface: Memory accesses in the range of A_0000h to B_FFFFh and IO address where address bits 9:0 are in the ranges of 3B0h to 3BBh or 3C0h to 3DFh. For IO cycles the decoding of address bits 15:10 depends on Vga16En.
18	IsaEn: ISA enable. Read-write.
17	SerrEn: SERR enable. Read-write.
16	ParityResponseEn: Parity response enable. Read-write. Controls the bridge's response to poisoned TLPs on its secondary interface. 1=The bridge takes its normal action when a poisoned TLP is received. 0=The bridge ignores any poisoned TLPs that it receives and continues normal operation.
15:11	IntPinR: interrupt pin. Read-only.
10:8	IntPin: interrupt pin. IF (D0F0xE4_x0[2:1]01_0010[HwInitWrLock]==1) THEN Read-only. ELSE Read-write. ENDIF.
7:0	IntLine: Interrupt line. Read-write.

D[8:2]F0x50 Power Management Capability Register

Reset: C803_5801h.

Bits	Description
31:27	PmeSupport. Read-only.
26	D2Support: D2 support. Read-only. D2 is not supported
25	D1Support: D1 support. Read-only. D1 is not supported
24:22	AuxCurrent: auxiliary current. Read-only. Auxiliary current is not supported.
21	DevSpecificInit: device specific initialization. Read-only. This field is hardwired to 0 to indicate that there is no device specific initialization necessary.
20	Reserved.
19	PmeClock. Read-only. 0=Indicate that PCI clock is not needed to generate PME messages.
18:16	Version: version. Read-only.
15:8	NextPtr: next pointer. Read-only. 58h=Address of the next capability structure.
7:0	CapID: capability ID. Read-only. 01h=PCI power management data structure.

D[8:2]F0x54 Power Management Control and Status Register

Bits	Description
31:24	PmeData. Read-only. Reset: 0.
23	BusPwrEn. Read-only. Reset: 0.
22	B2B3Support. Read-only. Reset: 0. B states are not supported.
21:16	Reserved.

15	PmeStatus: PME status. Read; Write-1-to-clear. Reset: 0. This bit is set when the root port would issue a PME message (independent of the state of the PmeEn bit). Once set, this bit remains set until it is reset by writing a 1 to this bit location. Writing a 0 has no effect.												
14:13	DataScale: data scale. Read-only. Reset: 0.												
12:9	DataSelect: data select. Read-only. Reset: 0.												
8	PmeEn: PME# enable. Read-write. Reset: 0.												
7:4	Reserved.												
3	NoSoftReset: no soft reset. Read-only. Reset: 0. Software is required to re-initialize the function when returning from D3 _{hot} .												
2	Reserved.												
1:0	PowerState: power state. Read-write. Reset: 0. This 2-bit field is used both to determine the current power state of the root port and to set the root port into a new power state. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>D0</td> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>01b</td> <td>Reserved</td> <td>11b</td> <td>D3</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	00b	D0	10b	Reserved	01b	Reserved	11b	D3
Bits	Definition	Bits	Definition										
00b	D0	10b	Reserved										
01b	Reserved	11b	D3										

D[8:2]F0x58 PCI Express Capability Register

Reset: 0042_A010h.

Bits	Description
31:30	Reserved.
29:25	IntMessageNum: interrupt message number. Read-only. This register indicates which MSI vector is used for the interrupt message.
24	SlotImplemented: Slot implemented. Read-only. 1=The IO Link associated with this port is connected to a slot.
23:20	DeviceType: device type. Read-only. 4h=Root complex.
19:16	Version. Read-only. 2h=GEN 2 compliant.
15:8	NextPtr: next pointer. Read-only. A0h=Pointer to the next capability structure.
7:0	CapID: capability ID. Read-only. 10h=PCIe [®] Capability structure.

D[8:2]F0x5C Device Capability Register

Reset: 0000_8020h.

Bits	Description
31:29	Reserved.
28	FlrCapable: function level reset capability. Read-only.
27:26	CapturedSlotPowerScale: captured slot power limit scale. Read-only.
25:18	CapturedSlotPowerLimit: captured slot power limit value. Read-only.
17:16	Reserved.
15	RoleBasedErrReporting: role-based error reporting. Read-only.
14:12	Reserved.

11:9	L1AcceptableLatency: endpoint L1 Acceptable Latency. Read-only.
8:6	L0SAcceptableLatency: endpoint L0s Acceptable Latency. Read-only.
5	ExtendedTag: extended tag support. Read-only. 1: 8 bit tag supported 0: 5 bit tag supported.
4:3	PhantomFunc: phantom function support. Read-only. 0=No phantom functions supported.
2:0	MaxPayloadSupport: maximum supported payload size. Read-only. 000b=128 bytes max payload size.

D[8:2]F0x60 Device Control and Status Register

Reset: 0000_2810h.

Bits	Description																
31:22	Reserved.																
21	TransactionsPending: transactions pending. Read-only. 0=No internally generated non-posted transactions pending.																
20	AuxPwr: auxiliary power. Read-only.																
19	UsrDetected: unsupported request detected. Read; Write-1-to-clear. 1=The port received an unsupported request. Errors are logged in this register even if error reporting is disabled.																
18	FatalErr: fatal error detected. Read; Write-1-to-clear. 1=The port detected a fatal error. Errors are logged in this register even if error reporting is disabled.																
17	NonFatalErr: non-fatal error detected. Read; Write-1-to-clear. 1=The port detected a non-fatal error. Errors are logged in this register even if error reporting is disabled.																
16	CorrErr: correctable error detected. Read; Write-1-to-clear. 1=The port detected a correctable error. Errors are logged in this register even if error reporting is disabled.																
15	BridgeCfgRetryEn: bridge configuration retry enable. Read-only.																
14:12	MaxRequestSize: maximum request size. Read-write.																
11	NoSnoopEnable: enable no snoop. Read-write. 1=The port is permitted to set the No Snoop bit in the Requester Attributes of transactions it initiates that do not require hardware enforced cache coherency.																
10	AuxPowerPmEn: auxiliary power PM enable. Read-only.																
9	PhantomFuncEn: phantom functions enable. Read-only.																
8	ExtendedTagEn: extended tag enable. Read-write. 1=8-bit tags generation enabled. 0=5-bit tags are used.																
7:5	MaxPayloadSize: maximum supported payload size. Read-write. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>128B</td> <td>3h</td> <td>1024B</td> </tr> <tr> <td>1h</td> <td>256B</td> <td>4h</td> <td>2048B</td> </tr> <tr> <td>2h</td> <td>512B</td> <td>5h</td> <td>4096B</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	0h	128B	3h	1024B	1h	256B	4h	2048B	2h	512B	5h	4096B
Bits	Definition	Bits	Definition														
0h	128B	3h	1024B														
1h	256B	4h	2048B														
2h	512B	5h	4096B														
4	RelaxedOrdEn: relaxed ordering enable. Read-write. 1=The root port is permitted to set the relaxed ordering bit in the attributes field of transactions it initiates that do not require strong write ordering.																

3	UsrReportEn: unsupported request reporting enable. Read-write. 1=Reporting of unsupported requests enabled.
2	FatalErrEn: fatal error reporting enable. Read-write. 1=Enable sending ERR_FATAL messages.
1	NonFatalErrEn: non-fatal error reporting enable. Read-write. 1=Enable sending ERR_NONFATAL messages.
0	CorrErrEn: correctable error reporting enable. Read-write. 1=Enable sending ERR_CORR messages.

D[8:2]F0x64 IO Link Capability Register

Bits	Description																		
31:24	PortNumber: port number. Value: 00h. This field indicates the port number for the given IO link.																		
23	Reserved.																		
22	AspmOptionalityCompliance. Value: 1.																		
21	LinkBWNotificationCap: link bandwidth notification capability. Value: 0.																		
20	DIActiveReportingCapable: data link layer active reporting capability. Value: 0.																		
19	SurpriseDownErrReporting. Value: 0.																		
18	ClockPowerManagement: clock power management. Value: 0. 0=Indicates that the reference clock must not be removed while in L1 or L2/L3 ready link states.																		
17:15	L1ExitLatency: L1 exit latency. Value: 010b. 010b=Indicate an exit latency between 2 us and 4 us.																		
14:12	L0sExitLatency: L0s exit latency. Value: 001b. 001b=Indicates an exit latency between 64 ns and 128 ns.																		
11:10	PMSupport: active state power management support. Value: 11b. 11b=Indicates support of L0s and L1.																		
9:4	LinkWidth: maximum link width. Value: 10h. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Reserved.</td> </tr> <tr> <td>01h</td> <td>1 lanes</td> </tr> <tr> <td>02h</td> <td>2 lanes</td> </tr> <tr> <td>04h</td> <td>4 lanes</td> </tr> <tr> <td>08h</td> <td>8 lanes</td> </tr> <tr> <td>0Ch</td> <td>12 lanes</td> </tr> <tr> <td>10h</td> <td>16 lanes</td> </tr> <tr> <td>3Fh-11h</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Definition	00h	Reserved.	01h	1 lanes	02h	2 lanes	04h	4 lanes	08h	8 lanes	0Ch	12 lanes	10h	16 lanes	3Fh-11h	Reserved.
Bits	Definition																		
00h	Reserved.																		
01h	1 lanes																		
02h	2 lanes																		
04h	4 lanes																		
08h	8 lanes																		
0Ch	12 lanes																		
10h	16 lanes																		
3Fh-11h	Reserved.																		
3:0	LinkSpeed: link speed. Value: IF (D[8:2]F0xE4_xA4[LcGen2EnStrap]==0) THEN 1h ELSE 2h ENDIF. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved.</td> </tr> <tr> <td>1h</td> <td>2.5 Gb/s</td> </tr> <tr> <td>2h</td> <td>5.0 Gb/s</td> </tr> <tr> <td>Fh-3h</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Definition	0h	Reserved.	1h	2.5 Gb/s	2h	5.0 Gb/s	Fh-3h	Reserved.								
Bits	Definition																		
0h	Reserved.																		
1h	2.5 Gb/s																		
2h	5.0 Gb/s																		
Fh-3h	Reserved.																		

D[8:2]F0x68 IO Link Control and Status Register

Reset: 1001_0000h.

Bits	Description														
31	LinkAutonomousBWStatus: link autonomous bandwidth status. IF (D[8:2]F0x64[LinkBWNotificationCap]==0) THEN Read-only. ELSE Read-write. ENDIF.														
30	LinkBWManagementStatus: link bandwidth management status. IF (D[8:2]F0x64[LinkBWNotificationCap]==0) THEN Read-only. ELSE Read-write. ENDIF.														
29	DLActive: data link layer link active. Read-only. This bit indicates the status of the data link control and management state machine. 1=DL_Active state. 0=All other states.														
28	SlotClockCfg: slot clock configuration. Read-only. 1=The root port uses the same clock that the platform provides.														
27	LinkTraining: link training. Read-only. This read-only bit indicates that the physical layer link training state machine 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 link training state machine exits the configuration/recovery state.														
26	Reserved.														
25:20	NegotiatedLinkWidth: negotiated link width. Read-only. This field indicates the negotiated width of the given PCI Express link. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>01h</td> <td>1 lanes</td> </tr> <tr> <td>02h</td> <td>2 lanes</td> </tr> <tr> <td>04h</td> <td>4 lanes</td> </tr> <tr> <td>08h</td> <td>8 lanes</td> </tr> <tr> <td>0Ch</td> <td>12 lanes</td> </tr> <tr> <td>10h</td> <td>16 lanes</td> </tr> </tbody> </table>	Bits	Definition	01h	1 lanes	02h	2 lanes	04h	4 lanes	08h	8 lanes	0Ch	12 lanes	10h	16 lanes
Bits	Definition														
01h	1 lanes														
02h	2 lanes														
04h	4 lanes														
08h	8 lanes														
0Ch	12 lanes														
10h	16 lanes														
19:16	LinkSpeed: link speed. Read-only. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0001b</td> <td>2.5 Gb/s</td> </tr> <tr> <td>0010b</td> <td>5 Gb/s</td> </tr> </tbody> </table>	Bits	Definition	0001b	2.5 Gb/s	0010b	5 Gb/s								
Bits	Definition														
0001b	2.5 Gb/s														
0010b	5 Gb/s														
15:12	Reserved.														
11	LinkAutonomousBWIntEn: link autonomous bandwidth interrupt enable. Read-write. 1=Enables the generation of an interrupt to indicate that the Link AutonomousBWStatus bit has been set.														
10	LinkBWManagementEn: link bandwidth management interrupt enable. Read-write. 1=Enables the generation of an interrupt to indicate that the LinkBWManagementStatus has been set.														
9	HWAutonomousWidthDisable: hardware autonomous width disable. Read-write. 1=Disables hardware from changing the link width for reasons other than attempting to correct unreliable link operation by reducing link width.														
8	ClockPowerManagementEn: clock power management enable. Read-write.														
7	ExtendedSync: extended sync. Read-write. 1=Forces the transmission of additional ordered sets when exiting the L0s state and when in the recovery state.														

6	CommonClockCfg: common clock configuration. Read-write. 1=Indicates that the root port and the component at the opposite end of this IO link are operating with a distributed common reference clock. 0=Indicates that the root port and the component at the opposite end of this IO Link are operating with asynchronous reference clock.												
5	RetrainLink: retrain link. Read-write; cleared-when-done. 1=Initiate link retraining.												
4	LinkDis: link disable. Read-write. 1=Disable link. Writes to this bit are immediately reflected in the value read from the bit, regardless of actual link state.												
3	ReadCplBoundary: read completion boundary. Read-only. 0=64 byte read completion boundary.												
2	Reserved.												
1:0	PmControl: active state power management enable. Read-write. This field controls the level of ASPM supported on the given IO link. <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;"><u>Bits</u></th> <th style="text-align: left;"><u>Definition</u></th> <th style="text-align: left;"><u>Bits</u></th> <th style="text-align: left;"><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Disabled.</td> <td>10b</td> <td>L1 Entry Enabled.</td> </tr> <tr> <td>01b</td> <td>L0s Entry Enabled.</td> <td>11b</td> <td>L0s and L1 Entry Enabled.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	00b	Disabled.	10b	L1 Entry Enabled.	01b	L0s Entry Enabled.	11b	L0s and L1 Entry Enabled.
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>										
00b	Disabled.	10b	L1 Entry Enabled.										
01b	L0s Entry Enabled.	11b	L0s and L1 Entry Enabled.										

D[8:2]F0x6C Slot Capability Register

Reset: 0004_0000h.

Bits	Description												
31:19	PhysicalSlotNumber: physical slot number. Read-write. This field indicates the physical slot number attached to this port. This field is set to a value that assigns a slot number that is unique within the chassis, regardless of the form factor associated with the slot. This field must be initialized to 0 for ports connected to devices that are on the system board.												
18	NoCmdCplSupport: no command completed support. Read-write. 1 =Indicates that this slot does not generate software notification when an issued command is completed by the hot-plug controller.												
17	ElecMechIIPresent: electromechanical interlock present. Read-write. 0=Indicates that a electro-mechanical interlock is not implemented for this slot.												
16:15	SlotPwrLimitScale: slot power limit scale. Read-write. Specifies the scale used for the SlotPwrLimitValue. Range of Values: <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;"><u>Bits</u></th> <th style="text-align: left;"><u>Definition</u></th> <th style="text-align: left;"><u>Bits</u></th> <th style="text-align: left;"><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1.0</td> <td>10b</td> <td>0.01</td> </tr> <tr> <td>01b</td> <td>0.1</td> <td>11b</td> <td>0.001</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>	00b	1.0	10b	0.01	01b	0.1	11b	0.001
<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>										
00b	1.0	10b	0.01										
01b	0.1	11b	0.001										
14:7	SlotPwrLimitValue: slot power limit value. Read-write. In combination with the SlotPwrLimitScale value, specifies the upper limit on power supplied by slot. Power limit (in Watts) calculated by multiplying the value in this field by the value in the SlotPwrLimitScale field.												
6	HotplugCapable: hot-plug capability. Read-write. 1=Indicates that this slot is capable of supporting hot-plug operations.												
5	HotplugSurprise: hot-plug surprise. Read-write. 1=Indicates that an adapter present in this slot might be removed from the system without any prior notification.												
4	PwrIndicatorPresent: power indicator present. Read-write. 0=Indicates that a power indicator is not implemented for this slot.												
3	AttnIndicatorPresent: attention indicator present. Read-write. 0=Indicates that a attention indicator is not implemented for this slot.												
2	MrlSensorPresent: manual retention latch sensor present. Read-write. 0=Indicates that a manual retention latch sensor is not implemented for this slot.												

1	PwrControllerPresent: power controller present. Read-write. 0=A power controller is not implemented for this slot.
0	AttnButtonPresent: attention button present. Read-write. 0=An attention button is not implemented for this slot.

D[8:2]F0x70 Slot Control and Status Register

IF (D[8:2]F0x58[SlotImplemented]==0) THEN Reset: 0040_0000h. ELSE Reset: 0000_0000h. ENDIF.

Bits	Description
31:25	Reserved.
24	DIStateChanged: data link layer state change. Read; Write-1-to-clear. This bit is set when the value reported in the D[8:2]F0x60[DIActive] is changed. In response to a data link layer state changed event, software must read D[8:2]F0x60[DIActive] to determine if the link is active before initiating configuration cycles to the hot plugged device.
23	ElecMechIISts: electromechanical interlock status. Read-only.
22	PresenceDetectState: presence detect state. Read-only; Updated-by-hardware. This bit indicates the presence of an adapter in the slot based on the physical layer in-band presence detect mechanism. The in-band presence detect mechanism requires that power be applied to an adapter for its presence to be detected. 0=Slot empty. 1=Card present in slot. For root ports not connected to slots (D[8:2]F0x58[SlotImplemented]=0b), this bit returns always 1.
21	MrlSensorState. Read-only.
20	CmdCpl: command completed. Read-only.
19	PresenceDetectChanged: presence detect changes. Read; Write-1-to-clear. This bit is set when the value reported in PresenceDetectState is changed.
18	MrlSensorChanged. Read; Write-1-to-clear.
17	PwrFaultDetected. Read; Write-1-to-clear.
16	AttnButtonPressed: attention button pressed. Read-only.
15:13	Reserved.
12	DIStateChangedEn: data link layer state changed enable. Read-write. 1=Enables software notification when D[8:2]F0x60[DIActive] is changed.
11	ElecMechIICntl: electromechanical interlock control. Read-only.
10	PwrControllerCntl: power controller control. Read-only.
9:8	PwrIndicatorCntl: power indicator control. Read-only.
7:6	AttnIndicatorControl: attention indicator control. Read-only.
5	HotplugIntrEn: hot-plug interrupt enable. Read-only.
4	CmdCplIntrEn: command complete interrupt enable. Read-only.
3	PresenceDetectChangedEn: presence detect changed enable. Read-only.
2	MrlSensorChangedEn: manual retention latch sensor changed enable. Read-only.
1	PwrFaultDetectedEn: power fault detected enable. Read-only.
0	AttnButtonPressedEn: attention button pressed enable. Read-only.

D[8:2]F0x74 Root Complex Capability and Control Register

Reset: 0001_0000h.

Bits	Description
31:17	Reserved.
16	CrSoftVisibility: CRS software visibility. Read-only. 1=Indicates that the root port supports returning configuration request retry status (CRS) completion status to software.
15:5	Reserved.
4	CrSoftVisibilityEn: CRS software visibility enable. Read-write. 1=Enables the root port returning configuration request retry status (CRS) completion status to software.
3	PmIntEn: PME interrupt enable. Read-write. 1=Enables interrupt generation upon receipt of a PME message as reflected D[8:2]F0x78[PmeStatus]. A PME interrupt is also generated if D[8:2]F0x78[PmeStatus]=1 and this bit is set by software.
2	SerrOnFatalErrEn: system error on fatal error enable. Read-write. 1=Indicates that a system error should be generated if a fatal error (ERR_FATAL) is reported by any of the devices in the hierarchy associated with this root port, or by the root port itself.
1	SerrOnNonFatalErrEn: system error on non-fatal error enable. Read-write. 1=Indicates that a system error should be generated if a non-fatal error (ERR_NONFATAL) is reported by any of the devices in the hierarchy associated with this root port, or by the root port itself.
0	SerrOnCorrErrEn: system error on correctable error enable. Read-write. 1=Indicates that a system error should be generated if a correctable error (ERR_COR) is reported by any of the devices in the hierarchy associated with this root port, or by the root port itself.

D[8:2]F0x78 Root Complex Status Register

Reset: 0000_0000h.

Bits	Description
31:18	Reserved.
17	PmePending: PME pending. Read-only. This bit indicates that another PME is pending when PmeStatus is set. When PmeStatus is cleared by software; the PME is delivered by hardware by setting the PmeStatus bit again and updating the requestor ID appropriately. PmePending is cleared by hardware if no more PMEs are pending.
16	PmeStatus: pme status. Read; Write-1-to-clear. This bit indicates that PME was asserted by the requestor ID indicated in the PmeRequestorID field. Subsequent PMEs are kept pending until PmeStatus is cleared by writing a 1.
15:0	PmeRequestorId: pme requestor ID. Read-only. This field indicates the PCI requestor ID of the last PME requestor.

D[8:2]F0x7C Device Capability 2

Reset: 0000_001Fh.

Bits	Description
31:6	Reserved.

5	AriForwardingSupported . Read-only.
4	CplTimeoutDisSup: completion timeout disable supported . Read-only.
3:0	CplTimeoutRangeSup: completion timeout range supported . Read-only. Fh=Completion timeout range is 64s to 50us.

D[8:2]F0x80 Device Control and Status 2

Reset: 0000_0000h.

Bits	Description																																
31:6	Reserved.																																
5	AriForwardingEn . Read-only.																																
4	CplTimeoutDis: completion timeout disable . Read-write. 1=Completion timeout disabled.																																
3:0	CplTimeoutValue: completion timeout value . Read-write. BIOS: 6h. <table border="1"> <thead> <tr> <th>Bits</th> <th>Timeout Range</th> <th>Bits</th> <th>Timeout Range</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>50ms-50us</td> <td>9h</td> <td>900ms-260ms</td> </tr> <tr> <td>1h</td> <td>100us-50us</td> <td>Ah</td> <td>3.5s-1s</td> </tr> <tr> <td>2h</td> <td>10ms-1ms</td> <td>Ch-Bh</td> <td>Reserved</td> </tr> <tr> <td>4h-3h</td> <td>Reserved</td> <td>Dh</td> <td>13s-4s</td> </tr> <tr> <td>5h</td> <td>55ms-16ms</td> <td>Eh</td> <td>64s-4s</td> </tr> <tr> <td>6h</td> <td>210ms-65ms</td> <td>Fh</td> <td>Reserved</td> </tr> <tr> <td>8h-7h</td> <td>Reserved</td> <td></td> <td></td> </tr> </tbody> </table>	Bits	Timeout Range	Bits	Timeout Range	0h	50ms-50us	9h	900ms-260ms	1h	100us-50us	Ah	3.5s-1s	2h	10ms-1ms	Ch-Bh	Reserved	4h-3h	Reserved	Dh	13s-4s	5h	55ms-16ms	Eh	64s-4s	6h	210ms-65ms	Fh	Reserved	8h-7h	Reserved		
Bits	Timeout Range	Bits	Timeout Range																														
0h	50ms-50us	9h	900ms-260ms																														
1h	100us-50us	Ah	3.5s-1s																														
2h	10ms-1ms	Ch-Bh	Reserved																														
4h-3h	Reserved	Dh	13s-4s																														
5h	55ms-16ms	Eh	64s-4s																														
6h	210ms-65ms	Fh	Reserved																														
8h-7h	Reserved																																

D[8:2]F0x84 IO Link Capability 2

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D[8:2]F0x88 IO Link Control and Status 2

Bits	Description
31:17	Reserved.
16	CurDeemphasisLevel: current deemphasis level . Read-only. Reset: D[8:2]F0xE4_xA4[LcGen2EnStrap] . 1=-3.5 dB. 0=-6 dB
15:13	Reserved.
12	ComplianceDeemphasis: compliance deemphasis . Read-write. Reset: 0. This bit defines the compliance deemphasis level when EnterCompliance is set. Software should leave this field in its default state. 1=-3.5 dB. 0=-6 dB
11	ComplianceSOS: compliance SOS . Read-write. Reset: 0. 1=The device transmits skip ordered sets in between the modified compliance pattern.
10	EnterModCompliance: enter modified compliance . Read-write. Reset: 0. 1=The device transmits modified compliance pattern. Software should leave this field in its default state.

9:7	XmitMargin: transmit margin. Read-write. Reset: 0. This field controls the non-deemphasized voltage level at the transmitter pins. Software should leave this field in its default state.										
6	SelectableDeemphasis: selectable deemphasis. Read-only. Reset: D[8:2]F0xE4_xA4[LcGen2EnStrap] . 0=Selectable deemphasis is not supported. 1=Selectable deemphasis supported.										
5	HwAutonomousSpeedDisable: hardware autonomous speed disable. Read-write. Cold reset: 0. 1=Support for hardware changing the link speed for device specific reasons disabled.										
4	EnterCompliance: enter compliance. Read-write. Cold reset: 0. 1=Force the link to enter the compliance mode.										
3:0	TargetLinkSpeed: target link speed. Read-write. Cold reset: {00b, D[8:2]F0xE4_xA4[LcGen2EnStrap] , ~ D[8:2]F0xE4_xA4[LcGen2EnStrap] }. This field defines the upper limit of the link operational speed. Writes of reserved encodings are not valid. Hardware prevents writes of reserved encodings from changing the state of this field. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved</td> </tr> <tr> <td>1h</td> <td>2.5GT/s</td> </tr> <tr> <td>2h</td> <td>5.0GT/s</td> </tr> <tr> <td>Fh-3h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	0h	Reserved	1h	2.5GT/s	2h	5.0GT/s	Fh-3h	Reserved
<u>Bits</u>	<u>Definition</u>										
0h	Reserved										
1h	2.5GT/s										
2h	5.0GT/s										
Fh-3h	Reserved										

D[8:2]F0x8C Slot Capability 2

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D[8:2]F0x90 Slot Control and Status 2

Reset: 0000_0000h.

Bits	Description
31:0	Reserved.

D[8:2]F0xA0 MSI Capability Register

Reset: 0000_B005h.

Bits	Description
31:24	Reserved.
23	Msi64bit: MSI 64 bit capability. Read-only. Value: D0F0x64_x46[Msi64bitEn] . 1=The device is capable of sending 64-bit MSI messages. 0=The device is not capable of sending a 64-bit message address.
22:20	MsiMultiEn: MSI multiple message enable. Read-write. Software writes to this field to indicate the number of allocated vectors (equal to or less than the number of requested vectors). When MSI is enabled, a function is allocated at least 1 vector.
19:17	MsiMultiCap: MSI multiple message capability. Read-only. 000b=The device is requesting one vector.

16	MsiEn: MSI enable. Read-write. 1=MSI generation is enabled and INTx generation is disabled. 0=MSI generation disabled and INTx generation is enabled.
15:8	NextPtr: next pointer. Read-only.
7:0	CapID: capability ID. Read-only. 05h=MSI capability structure.

D[8:2]F0xA4 MSI Message Address Low

Reset: 0000_0000h.

Bits	Description
31:2	MsiMsgAddrLo: MSI message address. Read-write. This register specifies the dword aligned address for the MSI memory write transaction.
1:0	Reserved.

IF (DOF0x64_x46[Msi64bitEn]==0) THEN

D[8:2]F0xA8 MSI Message Data

Reset: 0000_0000h.

Bits	Description
31:16	Reserved.
15:0	MsiData: MSI message data. Read-write. This register specifies lower 16 bits of data for the MSI memory write transaction. The upper 16 bits are always 0.

ELSE

D[8:2]F0xA8 MSI Message Address High

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7:0	MsiMsgAddrHi: MSI message address. Read-write. This register specifies the upper 8-bits of the MSI address in 64 bit MSI mode.

D[8:2]F0xAC MSI Message Data

Reset: 0000_0000h.

Bits	Description
31:16	Reserved.
15:0	MsiData: MSI message data. Read-write. This register specifies lower 16 bits of data for the MSI memory write transaction. The upper 16 bits are always 0.

ENDIF.

D[8:2]F0xB0 Subsystem and Subvendor Capability ID Register

Reset: 0000_B80Dh.

Bits	Description
31:16	Reserved.
15:8	NextPtr: next pointer. Read-only.
7:0	CapID: capability ID. Read-only.

D[8:2]F0xB4 Subsystem and Subvendor ID Register

Bits	Description
31:16	SubsystemID. Read-only. Value: D0F0xE4_x013[1:0]_0046[SubsystemID].
15:0	SubsystemVendorID. Read-only. Value: D0F0xE4_x013[1:0]_0046[SubsystemVendorID].

D[8:2]F0xB8 MSI Capability Mapping

Reset: A803_0008h.

Bits	Description
31:27	CapType: capability type. Read-only.
26:18	Reserved.
17	FixD. Read-only.
16	En. Read-only.
15:8	NextPtr: next pointer. Read-only.
7:0	CapID: capability ID. Read-only.

D[8:2]F0xBC MSI Mapping Address Low

Bits	Description
31:20	MsiMapAddrLo. Read-only. Reset: 0. Lower 32-bits of the MSI address.
19:0	Reserved.

D[8:2]F0xC0 MSI Mapping Address High

Bits	Description
31:0	MsiMapAddrHi. Read-only. Reset: 0. Upper 32-bits of the MSI address.

D[8:2]F0xE0 Root Port Index

Reset: 0000_0000h.

The index/data pair registers [D\[8:2\]F0xE0](#) and [D\[8:2\]F0xE4](#) is used to access the registers

D[8:2]F0xE4_x[FF:00]. To read or write to one of these registers, the address is written first into the address register D[8:2]F0xE0 and then the data are read or written by read or write the data register D[8:2]F0xE4.

Bits	Description
31:8	Reserved.
7:0	PcieIndex . Read-write.

D[8:2]F0xE4 Root Port Data

See D[8:2]F0xE0.

D[8:2]F0xE4_x02 Root Port Hardware Debug

Reset: 0000_0000h.

Bits	Description
31:16	Reserved.
15	RegsLcAllowTxL1Control . Read-write. BIOS: 1. 1=Tx can prevent LC from entering L1 when there are outstanding completions.
14:0	Reserved.

D[8:2]F0xE4_x20 Root Port TX Control

Reset: 0050_8000h.

Bits	Description
31:16	Reserved.
15	TxFlushTlpDis: TLP flush disable . Read-write. BIOS: ~D[8:2]F0x6C[HotplugCapable]. 1=Disable flushing TLPs when the link is down.
14:0	Reserved.

D[8:2]F0xE4_x50 Root Port Lane Status

Reset: 0000_0000h.

Bits	Description																				
31:7	Reserved.																				
6:1	PhyLinkWidth: port link width . Read-only. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00_0000b</td> <td>disabled</td> <td>00_1000b</td> <td>x8</td> </tr> <tr> <td>00_0001b</td> <td>x1</td> <td>01_0000b</td> <td>x12</td> </tr> <tr> <td>00_0010b</td> <td>x2</td> <td>10_0000b</td> <td>x16</td> </tr> <tr> <td>00_0100b</td> <td>x4</td> <td></td> <td></td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	00_0000b	disabled	00_1000b	x8	00_0001b	x1	01_0000b	x12	00_0010b	x2	10_0000b	x16	00_0100b	x4		
Bits	Definition	Bits	Definition																		
00_0000b	disabled	00_1000b	x8																		
00_0001b	x1	01_0000b	x12																		
00_0010b	x2	10_0000b	x16																		
00_0100b	x4																				
0	PortLaneReversal: port lane reversal . Read-only. 1=Port lanes order is reversed.																				

D[8:2]F0xE4_x6A Root Port Error Control

Reset: 0000_0500h.

Bits	Description
31:1	Reserved.
0	ErrReportingDis: advanced error reporting disable. Read-write. BIOS: 0. 1=Error reporting disabled. 0=Error reporting enabled.

D[8:2]F0xE4_x70 Root Port Receiver Control

Reset: 0000_43F7h.

Bits	Description																				
31:20	Reserved.																				
19	RxRcbCplTimeoutMode: RCB completion timeout mode. Read-write. BIOS: 1.																				
18:16	RxRcbCplTimeout: RCB completion timeout. Read-write. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Disabled</td> <td>100b</td> <td>50ms</td> </tr> <tr> <td>001b</td> <td>50us</td> <td>101b</td> <td>100ms</td> </tr> <tr> <td>010b</td> <td>10ms</td> <td>110b</td> <td>500ms</td> </tr> <tr> <td>011b</td> <td>25ms</td> <td>111b</td> <td>1ms</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	000b	Disabled	100b	50ms	001b	50us	101b	100ms	010b	10ms	110b	500ms	011b	25ms	111b	1ms
Bits	Definition	Bits	Definition																		
000b	Disabled	100b	50ms																		
001b	50us	101b	100ms																		
010b	10ms	110b	500ms																		
011b	25ms	111b	1ms																		
15:0	Reserved.																				

D[8:2]F0xE4_xA0 Per Port Link Controller (LC) Control

Reset: 4000_0050h.

Bits	Description																																				
31:24	Reserved.																																				
23	LcL1ImmediateAck: immediate ACK ASPM L1 entry. Read-write. BIOS: 1. 1=Always ACK ASPM L1 entry DLLPs.																																				
22:16	Reserved.																																				
15:12	LcL1Inactivity: L1 inactivity timer. Read-write. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>L1 disabled</td> <td>8h</td> <td>400us</td> </tr> <tr> <td>1h</td> <td>1us</td> <td>9h</td> <td>1ms</td> </tr> <tr> <td>2h</td> <td>2us</td> <td>Ah</td> <td>40us</td> </tr> <tr> <td>3h</td> <td>4us</td> <td>Bh</td> <td>10ms</td> </tr> <tr> <td>4h</td> <td>10us</td> <td>Ch</td> <td>40ms</td> </tr> <tr> <td>5h</td> <td>20us</td> <td>Dh</td> <td>100ms</td> </tr> <tr> <td>6h</td> <td>40us</td> <td>Eh</td> <td>400ms</td> </tr> <tr> <td>7h</td> <td>100us</td> <td>Fh</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	0h	L1 disabled	8h	400us	1h	1us	9h	1ms	2h	2us	Ah	40us	3h	4us	Bh	10ms	4h	10us	Ch	40ms	5h	20us	Dh	100ms	6h	40us	Eh	400ms	7h	100us	Fh	Reserved
Bits	Definition	Bits	Definition																																		
0h	L1 disabled	8h	400us																																		
1h	1us	9h	1ms																																		
2h	2us	Ah	40us																																		
3h	4us	Bh	10ms																																		
4h	10us	Ch	40ms																																		
5h	20us	Dh	100ms																																		
6h	40us	Eh	400ms																																		
7h	100us	Fh	Reserved																																		

11:8	LcL0sInactivity: L0s inactivity timer. Read-write.			
	<u>Bits</u>	<u>Definition</u>	<u>Bits</u>	<u>Definition</u>
	0h	L0s disabled	8h	4us
	1h	40ns	9h	10us
	2h	80ns	Ah	40us
	3h	120ns	Bh	100us
	4h	200ns	Ch	400us
	5h	400ns	Dh	1ms
	6h	1us	Eh	4ms
	7h	2us	Fh	Reserved
7:4	Lc16xClearTxPipe. Read-write. BIOS: 3h. Specifies the number of clock to drain the TX pipe.			
3:0	Reserved.			

D[8:2]F0xE4_xA1 LC Training Control

Reset: 0400_1080h.

Bits	Description
31:13	Reserved.
12	LcInitSpdChgWithCsrEn: enable software initialed speed change. Read-write. 1=Enable link speed negotiation when D[8:2]F0x68[RetrainLink]=1 and the target link speed differs from the current link speed. 0=Link speed negotiation will not be performed when D[8:2]F0x68[RetrainLink]=1.
11	LcDontGotoL0sifL1Armed: prevent Ls0 entry if L1 request in progress. Read-write. BIOS: 1. 1=Prevent the LTSSM from transitioning to Rcv_L0s if an acknowledged request to enter L1 is in progress.
10:0	Reserved.

D[8:2]F0xE4_xA2 LC Link Width Control

Reset: 0020_0006h.

Bits	Description												
31:23	Reserved.												
22:21	LcDynLanesPwrState: unused link power state. Read-write. Controls the state of unused links after a reconfiguration. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>on</td> <td>10b</td> <td>SB2</td> </tr> <tr> <td>01b</td> <td>SB1</td> <td>11b</td> <td>Off</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	00b	on	10b	SB2	01b	SB1	11b	Off
Bits	Definition	Bits	Definition										
00b	on	10b	SB2										
01b	SB1	11b	Off										
20	LcUpconfigCapable: upconfigure capable. Read-only. 1=Both ends of the link are upconfigure capable. 0=Both ends of the link are not upconfigure capable.												
19:14	Reserved.												
13	LcUpconfigureDis: upconfigure disable. Read-write. 1=Disable link upconfigure.												
12	LcUpconfigureSupport: upconfigure support. Read-write.												
11	LcShortReconfigEn: short re-configuration enable. Read-write. 1=Enable short link re-configuration												
10	LcRenegotiateEn: link reconfiguration enable. Read-write. 1=Enable link re-negotiation.												

9	LcRenegotiationSupport: re-negotiation support. Read-only. 1=Link re-negotiation is supported by the downstream device.																				
8	LcReconfigNow: re-configure link. Read-write; cleared-when-done. 1=Initiate link width change.																				
7	LcReconfigArcMissingEscape. Read-write. 1=Expedite transition from Recovery.Idle to Detect during a long reconfiguration.																				
6:4	LcLinkWidthRd: current link width. Read-only. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0</td> <td>100b</td> <td>8</td> </tr> <tr> <td>001b</td> <td>1</td> <td>101b</td> <td>12</td> </tr> <tr> <td>010b</td> <td>2</td> <td>110b</td> <td>16</td> </tr> <tr> <td>011b</td> <td>4</td> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Definition	Bits	Definition	000b	0	100b	8	001b	1	101b	12	010b	2	110b	16	011b	4	111b	Reserved
Bits	Definition	Bits	Definition																		
000b	0	100b	8																		
001b	1	101b	12																		
010b	2	110b	16																		
011b	4	111b	Reserved																		
3	Reserved.																				
2:0	LcLinkWidth: link width required. Read-write. See: LcLinkWidthRd.																				

D[8:2]F0xE4_xA3 LC Number of FTS Control

Reset: 00FF_020Ch.

Bits	Description
31:10	Reserved.
9	LcXmitFtsBeforeRecovery: transmit FTS before recovery. Read-write. 1=Transmit FTS before recovery.
8:0	Reserved.

D[8:2]F0xE4_xA4 LC Link Speed Control

Reset: 0440_0100h.

Bits	Description
31:28	Reserved.
27	LcMultUpstreamAutoSpdChngEn: enable multiple automatic speed changes. Read-write. 1=Enable multiple automatic speed changes.
26:20	Reserved.
19	LcOtherSideSupportsGen2: downstream link supports gen2. Read-only. 1=The downstream link currently supports gen2.
18:13	Reserved.
12	LcSpeedChangeAttemptFailed: speed change attempt failed. Read-only; updated-by-hardware. 1=LcSpeedChangeAttemptsAllowed has been reached.
11:10	Reserved.
9	LcInitiateLinkSpeedChange: initiate link speed change. Read-write; cleared-when-done. 1=Initiate link speed negotiation.
8:7	Reserved.
6	LcForceDisSwSpeedChange: force disable software speed changes. Read-write. 1=Force the PCIe core to disable speed changes initiated by private registers.

5:1	Reserved.
0	LcGen2EnStrap: gen2 PCIe support enable. Read-write. 1=Gen2 PCIe support enabled. 0=Gen2 PCIe support disabled.

D[8:2]F0xE4_xA5 LC State 0

Cold reset: 0000_0000h.

Table 123: Link controller state encodings

Bits	Description	Bits	Description	Bits	Description
00h	s_Detect_Quiet.	12h	Rcv_L0_and_Tx_L0s.	24h	s_Rcvd_Loopback.
01h	s_Start_common_Mode.	13h	Rcv_L0_and_Tx_L0s_FTS.	25h	s_Rcvd_Loopback_Idle.
02h	s_Check_Common_Mode.	14h	Rcv_L0s_and_Tx_L0.	26h	s_Rcvd_Reset_Idle.
03h	s_Rcvr_Detect.	15h	Rcv_L0s_and_Tx_L0_Idle.	27h	s_Rcvd_Disable_Entry.
04h	s_No_Rcvr_Loop	16h	Rcv_L0s_and_Tx_L0s.	28h	s_Rcvd_Disable_Idle.
05h	s_Poll_Quiet.	17h	Rcv_L0s_and_Tx_L0s_FTS.	29h	s_Rcvd_Disable.
06h	s_Poll_Active.	18h	s_L1_Entry.	2Ah	s_Detect_Idle.
07h	s_Poll_Compliance.	19h	s_L1_Idle.	2Bh	s_L23_Wait.
08h	s_Poll_Config.	1Ah	s_L1_Wait	2Ch	Rcv_L0s_Skp_and_Tx_L0.
09h	s_Config_Step1.	1Bh	s_L1.	2Dh	Rcv_L0s_Skp_and_Tx_L0_Idle.
0Ah	s_Config_Step3.	1Ch	s_L23_Stall.	2Eh	Rcv_L0s_Skp_and_Tx_L0s.
0Bh	s_Config_Step5.	1Dh	s_L23_Entry.	2Fh	Rcv_L0s_Skp_and_Tx_L0_FTS.
0Ch	s_Config_Step2.	1Eh	s_L23_Entry.	30h	s_Config_Step2b.
0Dh	s_Config_Step4.	1Fh	s_L23_Ready.	31h	s_Recovery_Speed.
0Eh	s_Config_Step6.	20h	s_Recovery_lock.	32h	s_Poll_Compliance_Idle.
0Fh	s_Config_Idle.	21h	s_Recovery_Config.	33h	s_Rcvd_Loopback_Speed.
10h	Rcv_L0_and_Tx_L0.	22h	s_Recovery_Idle.	3Fh-34h	Reserved.
11h	Rcv_L0_and_Tx_L0_Idle.	23h	s_Training_Bit.		

Bits	Description
31:30	Reserved.
29:24	LcPrevState3: previous link state 3. Read-only. See: Table 123 .
23:22	Reserved.
21:16	LcPrevState2: previous link state 2. Read-only. See: Table 123 .
15:14	Reserved.
13:8	LcPrevState1: previous link state 1. Read-only. See: Table 123 .
7:6	Reserved.
5:0	LcCurrentState: current link state. Read-only. See: Table 123 .

D[8:2]F0xE4_xB1 LC Control 2

Reset: 8608_0280h.

Bits	Description
31:21	Reserved.
20	LcBlockElIdleInL0: block electrical idle in l0. Read-write. BIOS: 1. 1=Prevent electrical idle from causing the receiver to transition from L0 to L0s.
19	LcDeassertRxEnInL0s: deassert RX_EN in L0s. Read-write. 1=Turn off transmitters in L0s.
18:0	Reserved.

D[8:2]F0xE4_xB5 LC Control 3

Reset: 2850_5020h.

Bits	Description										
31	Reserved.										
30	LcGoToRecovery: go to recovery. Read-write. 1=Force link in the L0 state to transition to the recovery state.										
29:4	Reserved.										
3	LcRcvdDeemphasis: received deemphasis. Read-only. Deemphasis advertised by the downstream device. 1=3.5dB. 0=6dB.										
2:1	LcSelectDeemphasisCntl: deemphasis control. Read-write. Specifies the deemphasis used by the transmitter. <table border="1" data-bbox="267 1108 1047 1281"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Use deemphasis from LcSelectDeemphasis.</td> </tr> <tr> <td>01b</td> <td>Use deemphasis advertised by the downstream device.</td> </tr> <tr> <td>10b</td> <td>6dB</td> </tr> <tr> <td>11b</td> <td>3.5dB</td> </tr> </tbody> </table>	Bits	Definition	00b	Use deemphasis from LcSelectDeemphasis.	01b	Use deemphasis advertised by the downstream device.	10b	6dB	11b	3.5dB
Bits	Definition										
00b	Use deemphasis from LcSelectDeemphasis.										
01b	Use deemphasis advertised by the downstream device.										
10b	6dB										
11b	3.5dB										
0	LcSelectDeemphasis: downstream deemphasis. Read-write. Specifies the downstream deemphasis. 1=3.5dB. 0=6dB.										

D[8:2]F0xE4_xC0 LC Strap Override

Bits	Description
31:16	Reserved.
15	StrapAutoRcSpeedNegotiationDis: autonomous speed negotiation disable strap override. Read-write. Reset: 1.
14	Reserved.
13	StrapForceCompliance: force compliance strap override. Read-write. Reset: 0b.
12:0	Reserved.

D[8:2]F0xE4_xC1 Root Port Miscellaneous Strap Override

Reset: 0000_0000h.

Bits	Description
31:3	Reserved.
2	StrapExtendedFmtSupported: Extended Fmt Supported strap override. Read-write. BIOS: See 2.12.2 [IOMMU Initialization].
1	StrapE2EPrefixEn: E2E Prefix En strap override. Read-write. BIOS: See 2.12.2 [IOMMU Initialization].
0	StrapReverseLanes: reverse lanes strap override. Read-write.

D[8:2]F0x100 Vendor Specific Enhanced Capability Register

Bits	Description
31:20	NextPtr: next pointer. Read-only. IF (D0F0xE4_x0[2:1]01_00B0[StrapF0AerEn] == 1) THEN Reset: 150h. ELSE Reset: 000h. ENDIF.
19:16	CapVer: capability version. Read-only. Reset: 1h.
15:0	CapID: capability ID. Read-only. Reset: 000Bh.

D[8:2]F0x104 Vendor Specific Header Register

Reset: 0101_0001h.

Bits	Description
31:20	VsecLen: vendor specific enhanced capability structure length. Read-only. Defined the number of bytes of the entire vendor specific enhanced capability structure including the header.
19:16	VsecRev: vendor specific enhanced capability version. Read-only.
15:0	VsecID: vendor specific enhanced capability ID. Read-only.

D[8:2]F0x108 Vendor Specific 1 Register

Reset: 0000_0000h.

Bits	Description
31:0	Scratch: scratch. Read-write. This field does not control any hardware.

D[8:2]F0x10C Vendor Specific 2 Register

Reset: 0000_0000h.

Bits	Description
31:0	Scratch: scratch. Read-write. This field does not control any hardware.

D[8:2]F0x128 Virtual Channel 0 Resource Status Register

Reset: 0002_0000h.

Bits	Description
31:18	Reserved.
17	VcNegotiationPending: virtual channel negotiation pending. Read-only. 1=Virtual channel negotiation in progress. This bit must be 0 before the virtual channel can be used.
16	PortArbTableStatus: port arbitration table status. Read-only.
15:0	Reserved.

D[8:2]F0x150 Advanced Error Reporting Capability

Reset: 0001_0001h.

Bits	Description
31:20	NextPtr: next pointer. Read-only.
19:16	CapVer: capability version. Read-only.
15:0	CapID: capability ID. Read-only.

D[8:2]F0x154 Uncorrectable Error Status

Cold reset: 0000_0000h.

Bits	Description
31:26	Reserved.
25	TlpPrefixStatus: TLP prefix blocked status. Read; Write-1-to-clear.
24	AtomicOpEgressBlockedTLPStatus: atomic op egress blocked TLP status. Read; Write-1-to-clear.
23	McBlockedTLPStatus: MC blocked TLP status. Read; Write-1-to-clear.
22	UncorrInternalErrStatus: uncorrectable internal error status. Read; Write-1-to-clear.
21	AcsViolationStatus: access control service status. Read; Write-1-to-clear.
20	UnsuppReqErrStatus: unsupported request error status. Read; Write-1-to-clear. The header of the unsupported request is logged.
19	EcrcErrStatus: end-to-end CRC error status. Read; Write-1-to-clear.
18	MalTlpStatus: malformed TLP status. Read; Write-1-to-clear. The header of the malformed TLP is logged.
17	RcvOvflStatus: receiver overflow status. Read-only.
16	UnexpCplStatus: unexpected completion timeout status. Read; Write-1-to-clear. The header of the unexpected completion is logged.
15	CplAbortErrStatus: completer abort error status. Read; Write-1-to-clear.
14	CplTimeoutStatus: completion timeout status. Read; Write-1-to-clear.
13	FcErrStatus: flow control error status. Read-only.

12	PsnErrStatus: poisoned TLP status. Read; Write-1-to-clear. The header of the poisoned transaction layer packet is logged.
11:6	Reserved.
5	SurprdnErrStatus: surprise down error status. Read-only. 0=Detection and reporting of surprise down errors is not supported.
4	DlpErrStatus: data link protocol error status. Read; Write-1-to-clear.
3:0	Reserved.

D[8:2]F0x158 Uncorrectable Error Mask

Cold reset: 0000_0000h.

Bits	Description
31:26	Reserved.
25	TlpPrefixMask: TLP prefix blocked mask. Read-only.
24	AtomicOpEgressBlockedTLPMask: atomic op egress blocked TLP mask. Read-only.
23	McBlockedTLPMask: MC blocked TLP mask. Read-only.
22	UncorrInternalErrMask: uncorrectable internal error mask. Read-only.
21	AcsViolationMask: access control service mask. Read-only. 1=ACS violation errors are not reported.
20	UnsuppReqErrMask: unsupported request error mask. Read-write. 1=Unsupported request errors are not reported.
19	EcrcErrMask: end-to-end CRC error mask. Read-only.
18	MalTlpMask: malformed TLP mask. Read-write. 1=Malformed TLP errors are not reported.
17	RcvOvflMask: receiver overflow mask. Read-only.
16	UnexpCplMask: unexpected completion timeout mask. Read-write. 1=Unexpected completion errors are not reported.
15	CplAbortErrMask: completer abort error mask. Read-only.
14	CplTimeoutMask: completion timeout mask. Read-write. 1=Completion timeout errors are not reported.
13	FcErrMask: flow control error mask. Read-only.
12	PsnErrMask: poisoned TLP mask. Read-write. 1=Poisoned TLP errors are not reported.
11:6	Reserved.
5	SurprdnErrMask: surprise down error mask. Read-only.
4	DlpErrMask: data link protocol error mask. Read-write. 1=Data link protocol errors are not reported.
3:0	Reserved.

D[8:2]F0x15C Uncorrectable Error Severity

Cold reset: 0006_2030h.

Bits	Description
31:26	Reserved.
25	TlpPrefixSeverity: TLP prefix blocked severity. Read-only.
24	AtomicOpEgressBlockedTLPSeverity: atomic op egress blocked TLP severity. Read-only.
23	McBlockedTLPSeverity: MC blocked TLP severity. Read-only.
22	UncorrInternalErrSeverity: uncorrectable internal error severity. Read-only.
21	AcsViolationSeverity: access control service severity. Read-only. 1=Fatal error. 0=Non-fatal error.
20	UnsuppReqErrSeverity: unsupported request error severity. Read-write. 1=Fatal error. 0=Non-fatal error.
19	EcreErrSeverity: end-to-end CRC error severity. Read-only.
18	MalTlpSeverity: malformed TLP severity. Read-write. 1=Fatal error. 0=Non-fatal error.
17	RcvOvflSeverity: receiver overflow severity. Read-only.
16	UnexpCplSeverity: unexpected completion timeout severity. Read-write. 1=Fatal error. 0=Non-fatal error.
15	CplAbortErrSeverity: completer abort error severity. Read-only.
14	CplTimeoutSeverity: completion timeout severity. Read-write. 1=Fatal error. 0=Non-fatal error.
13	FcErrSeverity: flow control error severity. Read-only.
12	PsnErrSeverity: poisoned TLP severity. Read-write. 1=Fatal error. 0=Non-fatal error.
11:6	Reserved.
5	SurprdnErrSeverity: surprise down error severity. Read-only.
4	DlpErrSeverity: data link protocol error severity. Read-write. 1=Fatal error. 0=Non-fatal error.
3:0	Reserved.

D[8:2]F0x160 Correctable Error Status

Cold reset: 0000_0000h.

Bits	Description
31:14	Reserved.
13	AdvisoryNonfatalErrStatus: advisory non-fatal error status. Read; Write-1-to-clear. 1=A non-fatal unsupported request errors or a non-fatal unexpected completion errors occurred.
12	ReplayTimerTimeoutStatus: replay timer timeout status. Read; Write-1-to-clear.
11:9	Reserved.
8	ReplayNumRolloverStatus: replay. Read; Write-1-to-clear. 1=The same transaction layer packet has been replayed three times and has caused the link to re-train.
7	BadDlpStatus: bad data link layer packet status. Read; Write-1-to-clear. 1=A link CRC error was detected.
6	BadTlpStatus: bad transaction layer packet status. Read; Write-1-to-clear. 1=A bad non-duplicated sequence ID or a link CRC error was detected.

5:1	Reserved.
0	RcvErrStatus: receiver error status. Read-only. 1=An 8B10B or disparity error was detected.

D[8:2]F0x164 Correctable Error Mask

Cold reset: 0000_2000h.

Bits	Description
31:14	Reserved.
13	AdvisoryNonfatalErrMask: advisory non-fatal error mask. Read-write. 1=Error is not reported.
12	ReplayTimerTimeoutMask: replay timer timeout mask. Read-write. 1=Error is not reported.
11:9	Reserved.
8	ReplayNumRolloverMask: replay. Read-write. 1=Error is not reported.
7	BadDlIpMask: bad data link layer packet mask. Read-write. 1=Error is not reported.
6	BadTlpMask: bad transaction layer packet mask. Read-write. 1=Error is not reported.
5:1	Reserved.
0	RcvErrMask: receiver error mask. Read-only. 1=Error is not reported.

D[8:2]F0x168 Advanced Error Control

Cold reset: 0000_0000h.

Bits	Description
31:9	Reserved.
8	EcrcCheckEn: data link protocol error severity. Read-only. 0=Specifies that end-to-end CRC generation is not supported.
7	EcrcCheckCap: data link protocol error severity. Read-only. 0=Specifies that end-to-end CRC check is not supported.
6	EcrcGenEn: end-to-end CRC enable. Read-only. 0=Specifies that end-to-end CRC generation is not supported.
5	EcrcGenCap: end-to-end CRC capability. Read-only. 0=Specifies that end-to-end CRC generation is not supported.
4:0	FirstErrPtr: first error pointer. Read-only. The First Error Pointer identifies the bit position of the first error reported in the Uncorrectable Error Status register.

D[8:2]F0x16C Header Log DW0

Cold reset: 0000_0000h.

Bits	Description
31:0	TlpHdr: transaction layer packet header log. Read-only. Contains the header for a transaction layer packet corresponding to a detected error. The upper byte represents byte 0 of the header.

D[8:2]F0x170 Header Log DW1

Cold reset: 0000_0000h.

Bits	Description
31:0	TlpHdr: transaction layer packet header log. Read-only. Contains the header for a transaction layer packet corresponding to a detected error. The upper byte represents byte 4 of the header.

D[8:2]F0x174 Header Log DW2

Cold reset: 0000_0000h.

Bits	Description
31:0	TlpHdr: transaction layer packet header log. Read-only. Contains the header for a transaction layer packet corresponding to a detected error. The upper byte represents byte 8 of the header.

D[8:2]F0x178 Header Log DW3

Cold reset: 0000_0000h.

Bits	Description
31:0	TlpHdr: transaction layer packet header log. Read-only. Contains the header for a transaction layer packet corresponding to a detected error. The upper byte represents byte 12 of the header.

D[8:2]F0x17C Root Error Command

Reset: 0000_0000h.

Bits	Description
31:3	Reserved.
2	FatalErrRepEn: fatal error reporting enable. Read-write. 1=Enables the generation of an interrupt when a fatal error is reported by any of the devices in the hierarchy associated with this Root Port.
1	NonfatalErrRepEn: non-fatal error reporting enable. Read-write. 1=Enables generation of an interrupt when a non-fatal error is reported by any of the devices in the hierarchy associated with this Root Port.
0	CorrErrRepEn: correctable error reporting enable. Read-write. 1=Enables generation of an interrupt when a correctable error is reported by any of the devices in the hierarchy associated with this Root Port.

D[8:2]F0x180 Root Error Status

Cold reset: 0000_0000h.

Bits	Description
31:27	AdvErrIntMsgNum: advanced error interrupt message number. Read-only.
26:7	Reserved.
6	NFatalErrMsgRcvd: fatal error message received. Read; Write-1-to-clear. Set to 1 when one or more fatal uncorrectable error messages have been received.

5	NonFatalErrMsgRcvd: non-fatal error message received. Read; Write-1-to-clear. Set to 1 when one or more non-fatal uncorrectable error messages have been received.
4	FirstUncorrFatalRcvd: first uncorrectable fatal error message received. Read; Write-1-to-clear. Set to 1 when the first uncorrectable error message received is for a fatal error.
3	MultErrFatalNonfatalRcvd: ERR_FATAL/NONFATAL message received. Read; Write-1-to-clear. Set when either a fatal or a non-fatal error is received and ErrFatalNonfatalRcvd is already set.
2	ErrFatalNonfatalRcvd: ERR_FATAL/NONFATAL message received. Read; Write-1-to-clear. Set when either a fatal or a non-fatal error is received and this bit is not already set.
1	MultErrCorrRcvd: multiple ERR_COR messages received. Read; Write-1-to-clear. Set when a correctable error message is received and ErrCorrRcvd is already set.
0	ErrCorrRcvd: ERR_COR message received. Read; Write-1-to-clear. Set when a correctable error message is received and this bit is not already set.

D[8:2]F0x184 Error Source ID

Cold reset: 0000_0000h.

Bits	Description
31:16	ErrFatalNonfatalSrcID: ERR_FATAL/ERR_NONFATAL source identification. Read-only. Loaded with the requestor ID indicated in the received ERR_FATAL or ERR_NONFATAL message when D[8:2]F0x180[ErrFatalNonfatalRcvd] is not already set.
15:0	ErrCorrSrcID: ERR_COR source identification. Read-only. Loaded with the requestor ID indicated in the received ERR_COR message when D[8:2]F0x180[ErrCorrRcvd] is not already set.

3.8 Device 18h Function 0 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space].

D18F0x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1400h.
15:0	VendorID: vendor ID. Read-only. Value: 1022h.

D18F0x04 Status/Command

Bits	Description
31:16	Status. Read-only. Value: 0010h. Bit[20] is set to indicate the existence of a PCI-defined capability block.
15:0	Command. Read-only. Value: 0000h.

D18F0x08 Class Code/Revision ID

Bits	Description
31:8	ClassCode. Read-only. Value: 060000h. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only. Value: 00h.

D18F0x0C Header Type

Read-only. Value: 0080_0000h.

Bits	Description
31:0	HeaderTypeReg. These bits are fixed at their default values. The header type field indicates that there are multiple functions present in this device.

D18F0x34 Capabilities Pointer

Bits	Description
31:8	Reserved.
7:0	CapPtr: capabilities pointer. Read-only. Value: 00h.

D18F0x40 Routing Table

Reset: 0004_0201h.

Bits	Description
31:0	Reserved.

D18F0x60 Node ID

Bits	Description										
31:21	Reserved.										
20:16	CpuCnt[4:0]: CPU count bits[4:0] . Read-write. Reset: 0. CpuCnt specifies the number of cores to be enabled in the system (the boot core of all nodes plus those cores enabled through D18F0x1DC[CpuEn]). <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>1 core.</td> </tr> <tr> <td>02h-01h</td> <td><CpuCnt[4:0]>+1 cores.</td> </tr> <tr> <td>03h</td> <td>4 cores.</td> </tr> <tr> <td>3Fh-04h</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	00h	1 core.	02h-01h	<CpuCnt[4:0]>+1 cores.	03h	4 cores.	3Fh-04h	Reserved.
Bits	Description										
00h	1 core.										
02h-01h	<CpuCnt[4:0]>+1 cores.										
03h	4 cores.										
3Fh-04h	Reserved.										
15:0	Reserved.										

D18F0x64 Unit ID

Reset: 0000_00E0h.

Bits	Description
31:8	Reserved.
7:6	HbUnit: host bridge Unit ID . Read-only. Specifies the coherent link Unit ID of the host bridge used by the coherent fabric.
5:4	MctUnit: memory controller Unit ID . Read-only. Specifies the coherent link Unit ID of the memory controller.
3:0	Reserved.

D18F0x68 Link Transaction Control

Bits	Description										
31:23	Reserved.										
22:21	DsNpReqLmt: downstream non-posted request limit . Read-write. Reset: 00b. BIOS: 01b. This specifies the maximum number of downstream non-posted requests issued by core(s) which may be outstanding on the IO links attached to this node at one time. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No limit</td> </tr> <tr> <td>01b</td> <td>limited to 1</td> </tr> <tr> <td>10b</td> <td>limited to 4</td> </tr> <tr> <td>11b</td> <td>limited to 8</td> </tr> </tbody> </table>	Bits	Description	00b	No limit	01b	limited to 1	10b	limited to 4	11b	limited to 8
Bits	Description										
00b	No limit										
01b	limited to 1										
10b	limited to 4										
11b	limited to 8										
20	SeqIdSrcNodeEn: sequence ID source node enable . Read-write. Reset: 0. 1=The source node ID of requests is provided in the SeqID field of the corresponding downstream IO link request packets. This may be useful for debug applications, in order to match downstream packets with their originating node. For normal operation, this bit should be cleared. Correct ordering of requests between different nodes is not ensured when this bit is set. Semaphore sharing between differing nodes may not work properly in systems which are capable of processing IO requests with differing non-zero SeqIDs out of request order.										

19	ApicExtSpur: APIC extended spurious vector enable. Read-write. Reset: 0. This enables the extended APIC spurious vector functionality; it affects APICF0[Vector]. 0=The lower 4 bits of the spurious vector are read-only 1111b. 1=The lower 4 bits of the spurious vector are writable.
18	ApicExtId: APIC extended ID enable. Read-write. Reset: 0. Enables the extended APIC ID functionality. 0=APIC ID is 4 bits. 1=APIC ID is 8 bits.
17	ApicExtBrdCst: APIC extended broadcast enable. Read-write. Reset: 0. Enables the extended APIC broadcast functionality. 0=APIC broadcast is 0Fh. 1=APIC broadcast is FFh. If ApicExtBrdCst=1 then software must assert ApicExtId.
16	LintEn: local interrupt conversion enable. Read-write. Reset: 0. 1=Enables the conversion of broadcast ExtInt and NMI interrupt requests to LINT0 and LINT1 local interrupts, respectively, before delivering to the local APIC. This conversion only takes place if the local APIC is hardware enabled. LINT0 and LINT1 are controlled by APIC3[60:50]. 0=ExtInt/NMI interrupts delivered unchanged.
15	LimitCldtCfg: limit coherent link configuration space range. Read-write. Reset: 0. BIOS: 1.
14:12	Reserved.
11	RespPassPW: response PassPW. Read-write. Reset: 0. 1=The PassPW bit in all downstream link responses is set, regardless of the originating request packet. This technically breaks the PCI ordering rules but it is not expected to be an issue in the downstream direction. Setting this bit improves the latency of upstream requests by allowing the downstream responses to pass posted writes. 0=The PassPW bit in downstream responses is based on the RespPassPW bit of the original request.
10	DisFillP: disable fill probe. Read-write. Reset: 0. Controls probes for core-generated fills. 0=Probes issued for cache fills. 1=Probes not issued for cache fills. BIOS: 0. BIOS may set if single core.
9	DisRmtPMemC: disable remote probe memory cancel. Read-write. Reset: 0. 1=Only probed caches on the same node as the target memory controller may generate MemCancel coherent link packets. MemCancels are used to attempt to save DRAM and/or link bandwidth associated with the transfer of stale DRAM data. 0=Probes hitting dirty blocks may generate MemCancel packets, regardless of the location of the probed cache.
8	DisPMemC: disable probe memory cancel. Read-write. Reset: 0. Controls generation of MemCancel coherent link packets. MemCancels are used to attempt to save DRAM and/or coherent link bandwidth associated with the transfer of stale DRAM data. 0=Probes hitting dirty blocks of the core cache may generate MemCancel packets. 1=Probes may not generate MemCancel packets.
7	CPURdRspPassPW: CPU read response PassPW. Read-write. Reset: 0. 1=Read responses to core-generated reads are allowed to pass posted writes. 0=core responses do not pass posted writes. This bit is not expected to be set. This bit may only be set during the boot process.
6	CPUReqPassPW: CPU request PassPW. Read-write. Reset: 0. 1=Core-generated requests are allowed to pass posted writes. 0=Core requests do not pass posted writes. This bit is not expected to be set. This bit may only be set during the boot process.
5	Reserved.
4	DisMTS: disable memory controller target start. Read-write. Reset: 0. 1=Disables use of TgtStart. TgtStart is used to improve scheduling of back-to-back ordered transactions by indicating when the first transaction is received and ordered at the memory controller.
3:0	Reserved.

D18F0x6C Link Initialization Control

Bits	Description
31	Reserved.
30	RlsLnkFullTokCntImm: release upstream full token count immediately. Read-write. Cold reset: 0. BIOS: 1 after buffer counts have been programmed. 1=Apply buffer counts programmed in D18F0x90 and D18F0x94 immediately without requiring warm reset. Once this bit is set, additional changes to the buffer counts only take effect upon warm reset.
29	Reserved.
28	RlsIntFullTokCntImm: release internal full token count immediately. Read-write. Cold reset: 0. BIOS: 1 after buffer counts have been programmed. 1=Apply buffer counts programmed in D18F3x6C , D18F3x70 , D18F3x74 , D18F3x78 , D18F3x7C , D18F3x140 , D18F3x144 , D18F3x148 , D18F3x17C , and D18F3x1A0 immediately without requiring warm reset. Once this bit is set, additional changes to the buffer counts only take effect upon warm reset.
27	ApplyIsocModeEnNow. Read-write. Cold reset: 0. BIOS: 1 after RlsLnkFullTokCntImm and RlsIntFullTokCntImm have been set. 1=Apply the programmed value in D18F0x84 [IsocEn] immediately without requiring warm reset. This bit may only be set if RlsLnkFullTokCntImm and RlsIntFullTokCntImm are set and isochronous buffers have been allocated. IF (ApplyIsocModeEnNow) THEN (D18F3x148 [IsocPreqTok0] > 0).
26:11	Reserved.
10:9	BiosRstDet[2:1]: BIOS reset detect bits[2:1]. See: BiosRstDet[0].
8:7	Reserved.
6	InitDet: CPU initialization command detect. Read-write. Reset: 0. This bit may be used by software to distinguish between an INIT and a warm/cold reset by setting it to a 1 before an initialization event is generated. This bit is cleared by RESET_L but not by an INIT command.
5	BiosRstDet[0]: BIOS reset detect bit[0]. Read-write. Cold reset: 0. BiosRstDet[2:0] = {BiosRstDet[2:1], BiosRstDet[0]}. May be used to distinguish between a reset event generated by the BIOS versus a reset event generated for any other reason by setting one or more of the bits to a 1 before initiating a BIOS-generated reset event.
4	ColdRstDet: cold reset detect. Read-write. Cold reset: 0. This bit may be used to distinguish between a cold versus a warm reset event by setting the bit to a 1 before an initialization event is generated.
3:1	Reserved.
0	RouteTblDis: routing table disable. Read-write. Reset: 1. BIOS: 0.

D18F0x84 Link Control

Cold reset: 7711_0000h.

This register is derived from the link control register defined in the link specification.

Bits	Description
31:13	Reserved.
12	IsocEn: isochronous flow-control mode enable. Read-write. BIOS: 1. This bit is set to place the link into isochronous flow-control mode (IFCM), as defined by the link specification. 1=IFCM. 0=Normal flow-control mode. See D18F0x6C [ApplyIsocModeEnNow].

11:6	Reserved.
5	Reserved.
4	LinkFail: link failure. Read; set-by-hardware; write-1-to-clear. This bit is set high by the hardware if a sync flood is received by the link. See 2.15.1.8.1 [Common Diagnosis Information].
3:0	Reserved.

D18F0x90 Upstream Base Channel Buffer Count

D18F0x90 and D18F0x94 specify the *hard-allocated* flow-control buffer counts in each virtual channel; it also provides the *free buffers* that may be used by any of the virtual channels, as needed, or reallocated by BIOS to the hard-allocated buffer counts. Base channel buffers are specified in D18F0x90; isochronous buffer counts (if in IFCM) are specified in D18F0x94. For all fields that specify buffer counts in D18F0x90 and D18F0x94 the number of buffers allocated is 2 times the value of the field.

The reset value is determined by whether the settings are locked by LockBc.

IF (LockBc) THEN Cold reset: 0005_01C9h.

ELSE (~LockBc) THEN Reset: 0005_01C9h.

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use D18F0x6C[RlsLnkFullTokCntImm] for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

The free list buffers (specified by FreeData and FreeCmd) are used to optimize buffer usage. When a transaction is received, if a free-list buffer is available, it is used for storage instead of one of the hard allocated buffers; as a result, a buffer release (for one of the hard allocated buffers used by the incoming request) can occur immediately without waiting for the transaction to be routed beyond the flow-control buffers.

Buffer allocation rules:

- The total number of command buffers allocated in the base and isochronous registers of a link cannot exceed 32.
 - $D18F0x90[NpReqCmd] + D18F0x90[PReq] + D18F0x90[RspCmd] + D18F0x90[ProbeCmd] + D18F0x90[FreeCmd] + D18F0x94[IsocNpReqCmd] + D18F0x94[IsocPReq] + D18F0x94[IsocRspCmd] \leq 16$.
- The total number of data buffers allocated in the base and isochronous registers of a link cannot exceed 16.
 - $D18F0x90[NpReqData] + D18F0x90[RspData] + D18F0x90[PReq] + D18F0x90[FreeData] + D18F0x94[IsocPReq] + D18F0x94[IsocNpReqData] + D18F0x94[IsocRspData] \leq 8$.
- The total number of hard allocated command buffers (ProbeCmd, RspCmd, PReq, NpReqCmd, and D18F0x94[IsocRspCmd, IsocRspCmd, IsocPReq, and IsocNpReqCmd]) cannot exceed 32.
 - $D18F0x90[ProbeCmd] + D18F0x90[RspCmd] + D18F0x90[PReq] + D18F0x90[NpReqCmd] + D18F0x94[IsocRspCmd] + D18F0x94[IsocRspCmd] + D18F0x94[IsocPReq] + D18F0x94[IsocNpReqCmd] \leq 16$.
- BIOS must set up non-zero counts (and adjust the base channel counts accordingly) prior to enabling IFCM.
 - IF (IFCM) THEN $(D18F0x94[IsocNpReqCmd]) > 0$.
- If an IOMMU is present in the system, D18F0x94[IsocNpReqCmd] must be non-zero.
 - IF (IOMMU) THEN $(D18F0x94[IsocNpReqCmd]) > 0$.

Bits	Description
31	LockBc: lock buffer count register. Read-write. Cold reset: 0. BIOS: 1. 1=The buffer count registers, D18F0x90 and D18F0x94 are locked such that warm resets do not place the registers back to their default value. Setting this bit does not prevent the buffer counts from being updated after a warm reset based on the value of the buffer counts before the warm reset. 0=Upon warm reset, the buffer count registers return to their default value after the link initializes regardless of the value before the warm reset.
30:28	Reserved.
27:25	FreeData: free data buffer count. Read-write. BIOS: 0.
24:20	FreeCmd: free command buffer count. Read-write. BIOS: 0.
19:18	RspData: response data buffer count. Read-write. BIOS: 1.
17:16	NpReqData: non-posted request data buffer count. Read-write. BIOS: 1.
15:12	ProbeCmd: probe command buffer count. Read-write. BIOS: 0.
11:8	RspCmd: response command buffer count. Read-write. BIOS: 2.
7:5	PReq: posted request command and data buffer count. Read-write. Specifies the number of posted command and posted data buffers allocated. BIOS: 5.
4:0	NpReqCmd: non-posted request command buffer count. Read-write. BIOS: 8.

D18F0x94 Link Isochronous Channel Buffer Count

See [D18F0x90](#).

The cold or warm reset is determined by whether the link initializes and whether the settings are locked by LockBc.

IF ([D18F0x90](#)[LockBc]) THEN Cold reset: 0000_0000h.
ELSE Reset: 0000_0000h.

Bits	Description
31:29	Reserved.
28:27	IsocRspData: isochronous response data buffer count. Read-write. BIOS: 0.
26:25	IsocNpReqData: isochronous non-posted request data buffer count. Read-write. BIOS: 1.
24:22	IsocRspCmd: isochronous response command buffer count. Read-write. BIOS: 0.
21:19	IsocPReq: isochronous posted request command and data buffer count. Read-write. BIOS: 0. This specifies the number of isochronous posted command and posted data buffers allocated.
18:16	IsocNpReqCmd: isochronous non-posted request command buffer count. Read-write. BIOS: 1.
15:8	SecBusNum: secondary bus number. Read-write. Specifies the configuration-space bus number of the IO link. When configured as a coherent link, this register has no meaning. This field should match the corresponding D18F1x[EC:E0] [BusNumBase], unless D18F1x[EC:E0] [DevCmpEn]=1, in which case this field should be 00h).
7:0	Reserved.

D18F0x98 Link Type

Bits	Description
31:6	Reserved.
5	PciEligible . Read-only. Reset: 1.
4:3	Reserved.
2	Reserved. Reset: 1.
1	Reserved. Reset: 1.
0	Reserved. Reset: 1.

D18F0x9C Link Frequency Extension

Bits	Description
31:16	Reserved.
15:1	Reserved. Reset: 1Fh.
0	Reserved.

D18F0x110 Link Clumping Enable

Reset: 0000_0000h. This register specifies how UnitIDs of upstream non-posted requests may be clumped per the link specification. The processor does not clump requests that it generates in the downstream direction.

Bits	Description
31:2	ClumpEn . Read-write. Each bit of this register corresponds to a link UnitID number. E.g., bit 2 corresponds to UnitID 02h, etc. 1=The specified UnitID is ordered in the same group as the specified UnitID - 1. For example if this register is programmed to 0000_00C0h, then UnitIDs 7h, 6h, and 5h are all ordered as if they are part of the same UnitID. This is used to allow more than 32 tags to be assigned to a single stream for the purposes of ordering.
1	Reserved.
0	Reserved.

D18F0x16C Link Global Extended Control

Reset: 0074_003Ah.

Bits	Description
31:0	Reserved.

D18F0x170 Link Extended Control

Reset: 0000_0001h.

Bits	Description
31:0	Reserved.

D18F0x1A0 Link Initialization Status

Reset: 0000_0000h.

Bits	Description								
31	InitStatusValid. Read-only. 1=Indicates that the rest of the information in this register is valid.								
30:2	Reserved.								
1:0	InitComplete. Read-only. <table border="1" data-bbox="267 493 1193 632"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Internal northbridge link has not completed initialization.</td> </tr> <tr> <td>10b-01b</td> <td>Reserved.</td> </tr> <tr> <td>11b</td> <td>Internal northbridge link has completed initialization.</td> </tr> </tbody> </table>	Bits	Description	00b	Internal northbridge link has not completed initialization.	10b-01b	Reserved.	11b	Internal northbridge link has completed initialization.
Bits	Description								
00b	Internal northbridge link has not completed initialization.								
10b-01b	Reserved.								
11b	Internal northbridge link has completed initialization.								

D18F0x1DC Core Enable

Reset: 0000_0000h.

Bits	Description
31:8	Reserved.
7:1	CpuEn: core enable. Read-write. This field is used to enable each of the cores after a reset. 1=Enable the core to start fetching and executing code from the boot vector. [1]: Core 1 enable; ...; [N]: Core N enable. The most significant bit N is indicated by CpuCoreNum, as defined in section 2.4.3 [Processor Cores and Downcoring]. All bits greater than N are reserved.
0	Reserved.

3.9 Device 18h Function 1 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space].

D18F1x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1401h.
15:0	VendorID: vendor ID. Read-only. Value: 1022h.

D18F1x08 Class Code/Revision ID

Bits	Description
31:8	ClassCode. Read-only. Value: 060000h. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only. Value: 00h. Processor revision. 00h=A0.

D18F1x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:0	HeaderTypeReg. Read-only. These bits are fixed at their default values. The header type field indicates that there are multiple functions present in this device.

D18F1x[144:140,44:40] DRAM Base/Limit

The following sets of registers specify the DRAM address ranges:

Base Low/High Limit Low/High
 F1x040, F1x140 F1x044, F1x144

F1x0XX registers provide the low address bits and F1x1XX registers provide the high address bits. Transaction addresses that are within the specified base/limit range are routed to the DstNode. See 2.8.2 [NB Routing].

DRAM mapping rules:

- Transaction addresses are within the defined range if:
 $\{\text{DramBase}[47:24], 00_0000\text{h}\} \leq \text{address}[47:0] \leq \{\text{DramLimit}[47:24], \text{FF_FFFFh}\}.$
- DRAM regions must not overlap each other.
- Accesses to addresses that map to both DRAM, as specified by the DRAM base and limit registers (F1x[1, 0][7C:40]), and MMIO, as specified by D18F1x[1CC:180,BC:80], are routed to MMIO only.
- Programming of the DRAM address maps must be consistent with the Memory-Type Range Registers (MTRRs) and the top of memory registers, MSRC001_001A and MSRC001_001D. CPU accesses only hit within the DRAM address maps if the corresponding MTRR is of type DRAM. Accesses from IO links are routed based on the DRAM base and limit registers (F1x[1, 0][7C:40]) only.
- The appropriate RE or WE bit(s) must be set. When initializing a base/limit pair, the BIOS must write the [limit] register before either the RE or WE bit is set. When changing a base/limit pair that is already enabled, the BIOS should clear RE and WE before changing the address range.
- See 2.8.2.1.1 [DRAM and MMIO Memory Space].

Hoisting. When memory hoisting is enabled in a node via [D18F1xF0](#)[DramHtHoleValid, DramHoleValid], the corresponding BaseAddr/LimitAddr should be configured to account for the memory hoisted above the hole. See [2.9.7 \[Memory Hoisting\]](#).

D18F1x40 DRAM Base Low

IF ([D18F2x118](#)[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0000_0000h.

Bits	Description
31:16	DramBase[39:24]: DRAM base address register bits[39:24]. DramBase[47:24] = { D18F1x140 [DramBase[47:40]], D18F1x40 [DramBase[39:24]]}.
15:2	Reserved.
1	WE: write enable. 1=Writes to this address range are enabled.
0	RE: read enable. 1=Reads to this address range are enabled.

D18F1x140 DRAM Base High

Bits	Description
31:8	Reserved.
7:0	DramBase[47:40]: DRAM base address register bits[47:40]. See: D18F1x40 [DramBase[39:24]].

D18F1x44 DRAM Limit Low

IF ([D18F2x118](#)[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF.

Bits	Description
31:16	DramLimit[39:24]: DRAM limit address register bits[39:24]. Reset: 00FCFFh. DramLimit[47:24] = { D18F1x144 [DramLimit[47:40]], D18F1x44 [DramLimit[39:24]]}.
15:11	Reserved.
10:8	Reserved.
7:3	Reserved.
2:0	DstNode: destination Node ID. Reset: 000b. Specifies the node that a packet is routed to if it is within the address range.

D18F1x144 DRAM Limit High

Bits	Description
31:8	Reserved.
7:0	DramLimit[47:40]: DRAM limit address register bits[47:40]. Reset: 00h. See: D18F1x44 [DramLimit[39:24]].

D18F1x[1CC:180,BC:80] MMIO Base/Limit

These registers, The memory mapped IO base and limit registers [D18F1x\[1CC:180,BC:80\]](#) specify the map-

ping from memory addresses to the corresponding node and IO link for MMIO transactions. Address ranges are specified by 12 sets of base/limit registers.

<u>Range</u>	<u>Base Low,High</u>	<u>Limit Low,High</u>
0	F1x80, F1x180	F1x84, F1x180
1	F1x88, F1x184	F1x8C, F1x184
2	F1x90, F1x188	F1x94, F1x188
3	F1x98, F1x18C	F1x9C, F1x18C
4	F1xA0, F1x190	F1xA4, F1x190
5	F1xA8, F1x194	F1xAC, F1x194
6	F1xB0, F1x198	F1xB4, F1x198
7	F1xB8, F1x19C	F1xBC, F1x19C
8	F1x1A0, F1x1C0	F1x1A4, F1x1C0
9	F1x1A8, F1x1C4	F1x1AC, F1x1C4
10	F1x1B0, F1x1C8	F1x1B4, F1x1C8
11	F1x1B8, F1x1CC	F1x1BC, F1x1CC

Transaction addresses that are within the specified base/limit range are routed to the node specified by DstNode and the link specified by DstLink. See 2.8.2 [NB Routing].

.MMIO mapping rules:

- Transaction addresses are within the defined range if:
 $\{\text{MMIOBase}[47:16], 0000\text{h}\} \leq \text{address}[47:0] \leq \{\text{MMIOLimit}[47:16], \text{FFFFh}\}$.
- MMIO regions must not overlap each other.
- Accesses to addresses that map to both DRAM, as specified by the DRAM base and limit registers (see D18F1x[144:140,44:40]), and MMIO, as specified by the memory mapped IO base and limit registers (F1x[BC:80]), are routed to MMIO only.
- Programming of the MMIO address maps must be consistent with the Memory-Type Range Registers (MTRRs) and the top of memory registers, MSRC001_001A and MSRC001_001D. CPU accesses only hit within the MMIO address maps if the corresponding MTRR is of type IO. Accesses from IO links are routed based on D18F1x[1CC:180,BC:80].
- The appropriate RE or WE bit(s) must be set. When initializing a base/limit pair, the BIOS must write the limit register before either the RE or WE bit is set. When changing a base/limit pair that is already enabled, the BIOS should clear RE and WE before changing the address range.
- Scenarios in which the address space of multiple MMIO ranges target the same IO device is supported.
- See 2.8.2.1.1 [DRAM and MMIO Memory Space].

D18F1x[1B8,1B0,1A8,1A0,B8,B0,A8,A0,98,90,88,80] MMIO Base Low

Table 124: Register Mapping for D18F1x[1B8,1B0,1A8,1A0,B8,B0,A8,A0,98,90,88,80]

Register	Function
D18F1x80	Range 0
D18F1x88	Range 1
D18F1x90	Range 2
D18F1x98	Range 3
D18F1xA0	Range 4
D18F1xA8	Range 5
D18F1xB0	Range 6

Table 124: Register Mapping for D18F1x[1B8,1B0,1A8,1A0,B8,B0,A8,A0,98,90,88,80]

D18F1xB8	Range 7
D18F1x1A0	Range 8
D18F1x1A8	Range 9
D18F1x1B0	Range 10
D18F1x1B8	Range 11

Bits	Description
31:8	MMIOBase[39:16]: MMIO base address register bits[39:16]. Read-write. Reset: 0. MMIOBase[47:16] = {D18F1x[1CC:1C0,19C:180][MMIOBase[47:40]], MMIOBase[39:16]}.
7:4	Reserved.
3	Lock. Read-write. Reset: 0. 1=the memory mapped IO base and limit registers (D18F1x[1CC:180,BC:80]) are read-only (including this bit). WE or RE in this register must be set in order for this to take effect.
2	Reserved.
1	WE: write enable. Read-write. Reset: 0. 1=Writes to this address range are enabled.
0	RE: read enable. Read-write. Reset: 0. 1=Reads to this address range are enabled.

D18F1x[1BC,1B4,1AC,1A4,BC,B4,AC,A4,9C,94,8C,84] MMIO Limit Low**Table 125: Register Mapping for D18F1x[1BC,1B4,1AC,1A4,BC,B4,AC,A4,9C,94,8C,84]**

Register	Function
D18F1x84	Range 0
D18F1x8C	Range 1
D18F1x94	Range 2
D18F1x9C	Range 3
D18F1xA4	Range 4
D18F1xAC	Range 5
D18F1xB4	Range 6
D18F1xBC	Range 7
D18F1x1A4	Range 8
D18F1x1AC	Range 9
D18F1x1B4	Range 10
D18F1x1BC	Range 11

Bits	Description
31:8	MMIOLimit[39:16]: MMIO limit address register bits[39:16]. Read-write. Reset: 0. MMIOLimit[47:16] = {D18F1x[1CC:1C0,19C:180][MMIOLimit[47:40]], MMIOLimit[39:16]}.

7	<p>NP: non-posted. Read-write. Reset: 0. 1=CPU write requests to this MMIO range are passed through the non-posted channel. This may be used to force writes to be non-posted for MMIO regions which map to the legacy ISA/LPC bus, or in conjunction with D18F0x68 [Link Transaction Control][DsN-pReqLmt] in order to allow downstream CPU requests to be counted and thereby limited to a specified number. This latter use of the NP bit may be used to avoid loop deadlock scenarios in systems that implement a region in an IO device that reflects downstream accesses back upstream. See the link summary of deadlock scenarios for more information. 0=CPU writes to this MMIO range use the posted channel. This bit does not affect requests that come from IO links (the virtual channel of the request is specified by the IO request).</p> <p>If two MMIO ranges target the same IO device and the NP bit is set differently in both ranges, unexpected transaction ordering effects are possible. In particular, using PCI- and IO-link-defined producer-consumer semantics, if a producer (e.g., the processor) writes data using a non-posted MMIO range followed by a flag to a posted MMIO range, then it is possible for the device to see the flag updated before the data is updated.</p>										
6	<p>DstSubLink: destination sublink. Read-write. Reset: 0. When a link is unganged, this bit specifies the destination sublink of the link specified by the memory mapped IO base and limit registers F1x[BC:80][DstLink]. 0=The destination link is sublink 0. 1=The destination link is sublink 1. If the link is ganged, then this bit must be low.</p>										
5:4	<p>DstLink: destination link ID. Read-write. Reset: 0. For transactions within the this MMIO range, this field specifies the destination IO link number of the destination node.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Link 0</td> </tr> <tr> <td>01b</td> <td>Link 1</td> </tr> <tr> <td>10b</td> <td>Link 2</td> </tr> <tr> <td>11b</td> <td>Link 3</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Link 0	01b	Link 1	10b	Link 2	11b	Link 3
<u>Bits</u>	<u>Description</u>										
00b	Link 0										
01b	Link 1										
10b	Link 2										
11b	Link 3										
3	Reserved.										
2:0	<p>DstNode: destination node ID bits. Read-write. Reset: 0. For transactions within the this MMIO range, this field specifies the destination node ID.</p>										

D18F1x[1CC:1C0,19C:180] MMIO Base/Limit High

Table 126: Register Mapping for D18F1x[D18F1x[1CC:1C0,19C:180]]

Register	Function
D18F1x180	Range 0
D18F1x184	Range 1
D18F1x188	Range 2
D18F1x18C	Range 3
D18F1x190	Range 4
D18F1x194	Range 5
D18F1x198	Range 6
D18F1x19C	Range 7
D18F1x1C0	Range 8

Table 126: Register Mapping for D18F1x[D18F1x[1CC:1C0,19C:180]

D18F1x1C4	Range 9
D18F1x1C8	Range 10
D18F1x1CC	Range 11

Bits	Description
31:24	Reserved.
23:16	MMIOLimit[47:40]: MMIO limit address register bits[47:40]. See: D18F1x[1BC,1B4,1AC,1A4,BC,B4,AC,A4,9C,94,8C,84][MMIOLimit[39:16]] .
15:8	Reserved.
7:0	MMIOBase[47:40]: MMIO base address register bits[47:40]. See: D18F1x[1B8,1B0,1A8,1A0,B8,B0,A8,A0,98,90,88,80][MMIOBase[39:16]] .

D18F1x[DC:C0] IO-Space Base/Limit

The IO-space base and limit registers, [D18F1x\[DC:C0\]](#), specify the mapping from IO addresses to the corresponding node and IO link for transactions resulting from x86-defined IN and OUT instructions. IO address ranges are specified by 4 sets of base/limit registers. The first set is F1xC0 and F1xC4, the second set is F1xC8 and F1xCC, and so forth. Transaction addresses that are within the specified base/limit range are routed to the node specified by DstNode and the link specified by DstLink. See [2.8.2 \[NB Routing\]](#).

IO mapping rules:

- IO-space transaction addresses are within the defined range if:
 - {IOBase[24:12], 000h} <= address <= {IOLimit[24:12], FFFh} and as specified by the IE bit; or
 - if the address is in the range specified by the VE bits.
- IO regions must not overlap each other.
- The appropriate RE or WE bit(s) must be set.
- See [2.8.2.1.2 \[IO Space\]](#).

D18F1x[D8,D0,C8,C0] IO-Space Base

Table 127: Register Mapping for D18F1x[D8,D0,C8,C0]

Register	Function
D18F1xC0	Range 0
D18F1xC8	Range 1
D18F1xD0	Range 2
D18F1xD8	Range 3

Bits	Description
31:25	Reserved.
24:12	IOBase[24:12]: IO base address register bits[24:12]. Read-write. Reset: 0.
11:6	Reserved.

5	IE: ISA enable. Read-write. Reset: 0. 1=The IO-space address window is limited to the first 256 B of each 1 KB block specified; this only applies to the first 64 KB of IO space. 0=The PCI IO window is not limited in this way.
4	VE: VGA enable. Read-write. Reset: 0. 1=Include IO-space transactions targeting the VGA-compatible address space within the IO-space window of this base/limit pair. These include IO accesses in which address bits[9:0] range from 3B0h to 3BBh or 3C0h to 3DFh (address bits[15:10] are not decoded); this only applies to the first 64 KB of IO space; i.e., address bits[24:16] must be low). 0=IO-space transactions targeting VGA-compatible address ranges are not added to the IO-space window. This bit should only be set in one register. The MMIO range associated with the VGA enable bit in the PCI specification is NOT included in the VE bit definition; to map this range to an IO link, see D18F1xF4 [VGA Enable] . When D18F1xF4[VE] is set, the state of this bit is ignored.
3:2	Reserved.
1	WE: write enable. Read-write. Reset: 0. 1=Writes to this IO-space address range are enabled.
0	RE: read enable. Read-write. Reset: 0. 1=Reads to this IO-space address range are enabled.

D18F1x[DC,D4,CC,C4] IO-Space Limit

Table 128: Register Mapping for D18F1x[DC,D4,CC,C4]

Register	Function
D18F1xC4	Range 0
D18F1xCC	Range 1
D18F1xD4	Range 2
D18F1xDC	Range 3

Bits	Description										
31:25	Reserved.										
24:12	IOLimit[24:12]: IO limit address register bits[24:12]. Read-write. Reset: 0.										
11:7	Reserved.										
6	DstSubLink: destination sublink. Read-write. Reset: 0. When a link is ungangled, this bit specifies the destination sublink of the link specified by F1x[DC:C0][DstLink]. 0=The destination link is sublink 0. 1=The destination link is sublink 1. If the link is gangled, then this bit must be low.										
5:4	DstLink: destination link ID. Read-write. Reset: 0. For transactions within the this IO-space range, this field specifies the destination IO link number of the destination node. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Link 0</td> </tr> <tr> <td>01b</td> <td>Link 1</td> </tr> <tr> <td>10b</td> <td>Link 2</td> </tr> <tr> <td>11b</td> <td>Link 3</td> </tr> </tbody> </table>	Bits	Description	00b	Link 0	01b	Link 1	10b	Link 2	11b	Link 3
Bits	Description										
00b	Link 0										
01b	Link 1										
10b	Link 2										
11b	Link 3										
3	Reserved.										
2:0	DstNode: destination node ID bits. Read-write. Reset: 0. For transactions within the this IO-space range, this field specifies the destination node ID.										

D18F1x[EC:E0] Configuration Map

D18F1x[EC:E0] specify the mapping from configuration address to the corresponding node and IO link. Configuration address ranges are specified by 4 pairs of base/limit registers. Transaction addresses that are within the specified base/limit range are routed to the node specified by DstNode and the link specified by DstLink. See 2.8.2 [NB Routing].

Table 129: Register Mapping for D18F1x[EC:E0]

Register	Function
D18F1xE0	Range 0
D18F1xE4	Range 1
D18F1xE8	Range 2
D18F1xEC	Range 3

Configuration space mapping rules:

- Configuration addresses (to “BusNo” and “Device” as specified by IOCF8 [IO-Space Configuration Address] in the case of IO accesses or 2.7 [Configuration Space] in the case of MMIO accesses) are within the defined range if:
 $(\{BusNumBase[7:0]\} \leq BusNo \leq \{BusNumLimit[7:0]\}) \& (DevCmpEn==0)$; or
 $(\{BusNumBase[4:0]\} \leq Device \leq \{BusNumLimit[4:0]\}) \& (DevCmpEn==1) \& (BusNo == 00h)$.
- Configuration regions must not overlap each other.
- The appropriate RE or WE bit(s) must be set.
- See 2.8.2.1.3 [Configuration Space].

Bits	Description
31:24	BusNumLimit[7:0]: bus number limit bits[7:0] . Read-write. Reset: 0.
23:16	BusNumBase[7:0]: bus number base bits[7:0] . Read-write. Reset: 0.
15:3	Reserved.
2	DevCmpEn: device number compare mode enable . Read-write. Reset: 0. 1=A device number range rather than a bus number range is used to specify the configuration-space window (see above).
1	WE: write enable . Read-write. Reset: 0. 1=Writes to this configuration-space address range are enabled.
0	RE: read enable . Read-write. Reset: 0. 1=Reads to this configuration-space address range are enabled.

D18F1xF0 DRAM Hole Address

IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. See 2.9.7 [Memory Hoisting].

Bits	Description
31:24	DramHoleBase[31:24]: DRAM hole base address . Reset: 0. Specifies the base address of the IO hole, below the 4GB address level, that is used in memory hoisting. Normally, DramHoleBase \geq MSRC001_001A[TOM[31:24]].
23:16	Reserved.
15:7	DramHoleOffset[31:23]: DRAM hole offset address . Reset: 0. When DramHoleValid=1, this offset is subtracted from the physical address of certain accesses in forming the normalized address.

6:3	Reserved.
2	Reserved.
1	DramMemHoistValid: dram memory hoist valid. Reset: 0. 1=Memory hoisting for the address range is enabled. 0=Memory hoisting is not enabled. This bit should be set if DramHoleValid=1.
0	DramHoleValid: dram hole valid. Reset: 0. 1=Memory hoisting is enabled. 0=Memory hoisting is not enabled. This bit should be set in the node(s) that own the DRAM address space that is hoisted above the 4 GB address level. See DramHoleOffset.

D18F1xF4 VGA Enable

Reset: 0000_0000h. All these bits are read-write unless Lock is set.

Bits	Description										
31:15	Reserved.										
14	DstSubLink: destination sublink. Read-write. When a link is ungangled, this bit specifies the destination sublink of the link specified by D18F1xF4[DstLink] . 0=The destination link is sublink 0. 1=The destination link is sublink 1. If the link is gangled, then this bit must be low.										
13:12	DstLink: destination link ID. Read-write. For transactions within the D18F1xF4[VE] -defined ranges, this field specifies the destination IO link number of the destination node. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Link 0</td> </tr> <tr> <td>01b</td> <td>Link 1</td> </tr> <tr> <td>10b</td> <td>Link 2</td> </tr> <tr> <td>11b</td> <td>Link 3</td> </tr> </tbody> </table>	Bits	Description	00b	Link 0	01b	Link 1	10b	Link 2	11b	Link 3
Bits	Description										
00b	Link 0										
01b	Link 1										
10b	Link 2										
11b	Link 3										
11:7	Reserved.										
6:4	DstNode: destination node ID. Read-write. For transactions within the D18F1xF4[VE] -defined range, this field specifies the destination node ID.										
3	Lock. Read-write. 1=All the bits in this register (D18F1xF4) are read-only (including this bit).										
2	Reserved.										
1	NP: non-posted. Read-write. 1=CPU write requests to the D18F1xF4[VE] -defined MMIO range are passed through the non-posted channel. 0=CPU writes may be posted.										
0	VE: VGA enable. Read-write. 1=Transactions targeting the VGA-compatible address space are routed and controlled as specified by this register. The VGA-compatible address space is: (1) the MMIO range A_0000h through B_FFFFh; (2) IO-space accesses in which address bits[9:0] range from 3B0h to 3BBh or 3C0h to 3DFh (address bits[15:10] are not decoded; this only applies to the first 64 KB of IO space; i.e., address bits[24:16] must be low). 0=Transactions targeting the VGA-compatible address space are not affected by the state of this register. When this bit is set, the state of D18F1x[DC:C0][VE] is ignored.										

D18F1x10C DCT Configuration Select

Reset: 0000_0000h.

Bits	Description
31:6	Unused.

5:4	<p>NbPsSel: NB P-state configuration select. Read-write. Read-write. Specifies the set of DCT Pstate registers to which accesses are routed.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>NB P-state 0</td> </tr> <tr> <td>01b</td> <td>NB P-state 1</td> </tr> <tr> <td>10b</td> <td>NB P-state 2</td> </tr> <tr> <td>11b</td> <td>NB P-state 3</td> </tr> </tbody> </table> <p>The following registers must be programmed for each NB P-state enabled by D18F5x1[6C:60][NbP-stateEn]:</p> <ul style="list-style-type: none"> D18F2x210_dct[1:0]_nbp[3:0][MaxRdLatency, DataTxFifoWrDly, RdPtrInit]. 	Bits	Description	00b	NB P-state 0	01b	NB P-state 1	10b	NB P-state 2	11b	NB P-state 3
Bits	Description										
00b	NB P-state 0										
01b	NB P-state 1										
10b	NB P-state 2										
11b	NB P-state 3										
3	<p>MemPsSel: Memory P-state configuration select. Read-write. Specifies the set of DCT controller registers to which accesses are routed. This register works independently of NbPsSel. 0=Memory P-state 0. 1=Memory P-state 1. See 2.5.7.1 [Memory P-states] and 2.9.1 [DCT Configuration Registers]. The following registers must be programmed for each memory P-state enabled by D18F5x1[6C:60][MemPstate]:</p> <ul style="list-style-type: none"> D18F2x2E8_dct[1:0]_mp[1:0] D18F2x2EC_dct[1:0]_mp[1:0] 										
2:1	Reserved.										
0	<p>DctCfgSel: DRAM controller configuration select. Read-write. Specifies DCT controller to which accesses are routed. 0=DCT 0. 1=DCT 1. See 2.9.1 [DCT Configuration Registers].</p>										

D18F1x120 DRAM Base System Address

IF ([D18F2x118\[LockDramCfg\]](#)) THEN Read-only. ELSE Read-write. ENDIF. Cold reset: 0000_0000h. [D18F1x120](#) and [D18F1x124](#) are required to specify the base and limit system address range of the DRAM connected to the local node. DRAM accesses to the local node with physical address [Addr\[47:0\]](#) that are within the following range are directed to the DCTs:

{[DramBaseAddr\[47:27\]](#), 000_0000h} <= [Addr\[47:0\]](#) <= {[DramLimitAddr\[47:27\]](#), 7FF_FFFFh};

DRAM accesses to the local node that are outside of this range are master aborted.

DRAM may be mapped as continuous regions for each node or it may be interleaved between nodes. If node interleaving is not invoked, as specified by [DramIntlvEn](#), then the address of the DRAM transaction is normalized before passing it to the DCTs by subtracting [DramBaseAddr](#).

Bits	Description
31:21	Reserved.
20:0	DramBaseAddr[47:27] .

D18F1x124 DRAM Limit System Address

IF ([D18F2x118\[LockDramCfg\]](#)) THEN Read-only. ELSE Read-write. ENDIF. See [D18F1x120 \[DRAM Base System Address\]](#).

Bits	Description
31:21	Reserved
20:0	DramLimitAddr[47:27] . Cold reset: 1F_FFFFh.

3.10 Device 18h Function 2 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space].

D18F2x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1402h.
15:0	VendorID: vendor ID. Read-only. Value: 1022h.

D18F2x08 Class Code/Revision ID

Reset: 0600_0000h.

Bits	Description
31:8	ClassCode. Read-only. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only.

D18F2x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:0	HeaderTypeReg. Read-only. These bits are fixed at their default values. The header type field indicates that there multiple functions present in this device.

D18F2x[5C:40]_dct[1:0] DRAM CS Base Address

IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers].

These registers along with D18F2x[6C:60]_dct[1:0] [DRAM CS Mask], translate DRAM request addresses (to a DRAM controller) into DRAM chip selects. Supported DIMM sizes are specified in D18F2x80_dct[1:0] [DRAM Bank Address Mapping]. For more information on the DRAM controllers, see 2.9 [DRAM Controllers (DCTs)].

For each chip select, there is a DRAM CS Base Address register. For every two chip selects there is a DRAM CS Mask Register.

Table 130: DIMM, Chip Select, and Register Mapping

Base Address Registers	Mask Register	Logical DIMM	Chip Select ¹
D18F2x40_dct[1:0]	F2x60	0	MEMCS[1:0]_L[0]
D18F2x44_dct[1:0]			MEMCS[1:0]_L[1]
D18F2x48_dct[1:0]	F2x64	1	MEMCS[1:0]_L[2]
D18F2x4C_dct[1:0]			MEMCS[1:0]_L[3]
D18F2x50_dct[1:0]	F2x68	Reserved	Reserved.
D18F2x54_dct[1:0]			Reserved.

Table 130: DIMM, Chip Select, and Register Mapping

Base Address Registers	Mask Register	Logical DIMM	Chip Select ¹
D18F2x58_dct[1:0]	F2x6C	Reserved	Reserved.
D18F2x5C_dct[1:0]			Reserved.
1. See 2.9.2 [DDR Pad to Processor Pin Mapping]			

The DRAM controller operates on the normalized physical address of the DRAM request. The normalized physical address includes all of the address bits that are supported by a DRAM controller. See 2.8 [Northbridge (NB)].

Each base address register specifies the starting normalized address of the block of memory associated with the chip select. Each mask register specifies the additional address bits that are consumed by the block of memory associated with the chip selects. If both chip selects of a DIMM are used, they must be the same size; in this case, a single mask register covers the address space consumed by both chip selects.

Lower-order address bits are provided in the base address and mask registers, as well. These allow memory to be interleaved between chip selects, such that contiguous physical addresses map to the same DRAM page of multiple chip selects. See 2.9.6.1 [Chip Select Interleaving]. The hardware supports the use of lower-order address bits to interleave chip selects if (1) the each chip select of the memory system spans the same amount of memory and (2) the number of chip selects of the memory system is a power of two.

System BIOS is required to assign the largest DIMM chip-select range to the lowest normalized address of the DRAM controller. As addresses increase, the chip-select size is required to remain constant or decrease. This is necessary to keep DIMM chip-select banks on aligned address boundaries, regardless as to the amount of address space covered by each chip select.

For each normalized address for requests that enters a DRAM controller, a ChipSelect[i] is asserted if:

```
CSEnable[i] &
( { (InputAddr[38:27] & ~AddrMask[i][38:27]),
    (InputAddr[21:11] & ~AddrMask[i][21:11]) } ==
  { (BaseAddr[i][38:27] & ~AddrMask[i][38:27]),
    (BaseAddr[i][21:11] & ~AddrMask[i][21:11]) } );
```

Bits	Description
31	Reserved.
30:19	BaseAddr[38:27]: normalized physical base address bits [38:27].
18:16	Reserved.
15:5	BaseAddr[21:11]: normalized physical base address bits [21:11].
4	Reserved.

3	OnDimmMirror: on-DIMM mirroring (ODM) enabled. 1=Address and bank bits are swapped by hardware for MRS commands sent to this chip select. This mode accounts for routing on the DIMM. Hardware bit swapping does not occur for commands sent via <code>D18F2x7C_dct[1:0][SendMrsCmd]</code> when <code>D18F2x7C_dct[1:0][EnDramInit] = 0</code> . This bit is expected to be set for the odd numbered rank of unbuffered DDR3 DIMMs if SPD byte 63 indicates that address mapping is mirrored. See 2.9.5.8.1.1 [DDR3 MR Initialization] . The following bits are swapped when enabled: <ul style="list-style-type: none"> • BA0 and BA1. • A3 and A4. • A5 and A6. • A7 and A8.
2	TestFail: memory test failed. Set by BIOS to indicate that a rank is present but has failed memory training or a memory consistency test, indicating that the memory is bad. BIOS should treat <code>CSEnable=1</code> and <code>TestFail=1</code> as mutually exclusive.
1	Reserved.
0	CSEnable: chip select enable.

D18F2x[6C:60]_dct[1:0] DRAM CS Mask

IF (`D18F2x118[LockDramCfg]`) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0000_0000h. See [2.9.1 \[DCT Configuration Registers\]](#). See `D18F2x[5C:40]_dct[1:0]`.

Bits	Description
31	Reserved.
30:19	AddrMask[38:27]: normalized physical address mask bits [38:27].
18:16	Reserved.
15:5	AddrMask[21:11]: normalized physical address mask bits [21:11].
4:0	Reserved.

D18F2x78_dct[1:0] DRAM Control

See [2.9.1 \[DCT Configuration Registers\]](#).

Bits	Description
31:18	Reserved.
17	AddrCmdTriEn: address command tristate enable. Read-write. Reset: 0. BIOS: See 2.9.5.7.1=Tristate the address, command, and bank buses when a Deselect command is issued.
16:0	Reserved.

D18F2x7C_dct[1:0] DRAM Initialization

Reset: 0000_0000h. See [2.9.1 \[DCT Configuration Registers\]](#). See [2.9.5.8.1 \[Software DDR3 Device Initialization\]](#).

BIOS can directly control the DRAM initialization sequence using this register. BIOS sets `EnDramInit` to start DRAM initialization. BIOS should then complete the initialization sequence specified in the appropriate JEDEC specification. After completing the sequence, BIOS clears `EnDramInit` to complete DRAM initializa-

tion. Setting more than one of the command bits in this register (SendControlWord, SendMrsCmd, and SendAutoRefresh) at a time results in undefined behavior.

Bits	Description								
31	EnDramInit: enable DRAM initialization. Read-write. 1=Place the DRAM controller in the BIOS-controlled DRAM initialization mode. The DCT asserts memory reset and deasserts CKE when this bit is set. BIOS must wait until <code>D18F2x98_dct[1:0][DctAccessDone] = 1</code> before programming <code>AssertCke=1</code> and <code>DeassertMemRstX=1</code> . BIOS must clear this bit after DRAM initialization is complete. BIOS must not set this bit on a DCT with no attached DIMMs. See 2.9.5.8.1 [Software DDR3 Device Initialization].								
30	SendControlWord: send control word. Read; write-1-only; cleared-by-hardware. 1= The DCT sends a control word to a chip select pair defined in <code>D18F2xA8_dct[1:0][CtrlWordCS]</code> . This bit is cleared by hardware after the command completes. This bit is valid only when <code>D18F2x90_dct[1:0][UnbuffDimm] = 0</code> .								
29	SendZQCmd: send ZQ command. Read; write-1-only; cleared-by-hardware. 1=The DCT sends the ZQ calibration command with either all even or all odd chip selects active. The first command targets even chip selects. Subsequent commands alternate between even and odd chip selects. This bit is cleared by the hardware after the command completes. This bit is valid only when <code>EnDramInit=1</code> .								
28	AssertCke: assert CKE. Read-write. Setting this bit causes the DCT to assert the CKE pins. This bit cannot be used to deassert the CKE pins.								
27	DeassertMemRstX: deassert memory reset. Read-write. Setting this bit causes the DCT to deassert the memory reset. This bit cannot be used to assert the memory reset pin.								
26	SendMrsCmd: send MRS command. Read; write-1-only; cleared-by-hardware. 1=The DCT sends the MRS commands defined by the <code>MrsAddress</code> and <code>MrsBank</code> fields of this register. This bit is cleared by hardware after the command completes. See 2.9.5.8.1.1 [DDR3 MR Initialization].								
25	SendAutoRefresh: send auto refresh command. Read; write-1-only; cleared-by-hardware. 1=The DCT sends an auto refresh command. This bit is cleared by hardware after the command completes.								
24	Reserved.								
23:21	MrsChipSel: MRS command chip select. Read-write. Specifies which DRAM chip select is used for MRS commands. Defined only if (<code>~EnDramInit ~D18F2x90_dct[1:0][UnbuffDimm]</code>); otherwise MRS commands are sent to all chip selects. <table border="1" data-bbox="267 1339 1039 1480"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>MRS command is sent to CS0</td> </tr> <tr> <td>110b-001b</td> <td>MRS command is sent to CS<MrsChipSel></td> </tr> <tr> <td>111b</td> <td>MRS command is sent to CS7</td> </tr> </tbody> </table>	Bits	Description	000b	MRS command is sent to CS0	110b-001b	MRS command is sent to CS<MrsChipSel>	111b	MRS command is sent to CS7
Bits	Description								
000b	MRS command is sent to CS0								
110b-001b	MRS command is sent to CS<MrsChipSel>								
111b	MRS command is sent to CS7								
20:18	MrsBank[2:0]: bank address for MRS commands. Read-write. Specifies the data driven on the DRAM bank pins for MRS commands.								
17:0	MrsAddress[17:0]: address for MRS commands. Read-write. Specifies the data driven on the DRAM address pins for MRS commands.								

D18F2x80_dct[1:0] DRAM Bank Address Mapping

IF (`D18F2x118[LockDramCfg]`) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. These fields specify DIMM configuration information. These fields are required to be programmed per the following table, based on the DRAM device size and with information of the DIMM. Table 131 show the bit numbers for each position.

Bits	Description
31:16	Reserved.
15:12	DimmAddrMap3: DIMM 3 address map.
11:8	DimmAddrMap2: DIMM 2 address map.
7:4	DimmAddrMap1: DIMM 1 address map.
3:0	DimmAddrMap0: DIMM 0 address map.

Table 131: DDR3 DRAM Address Mapping

Bits	CS Size	Device size, width	Bank			Address																	
			2	1	0		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
0000b		Reserved				Row																	
						Col																	
0001b	256MB	512Mb, x16	15	14	13	Row	x	x	x	x	17	16	27	26	25	24	23	22	21	20	19	18	
						Col	x	x	x	x	x	AP	12	11	10	9	8	7	6	5	4	3	
0010b	512MB	512Mb, x8 1Gb, x16	15	14	13	Row	x	x	x	17	16	28	27	26	25	24	23	22	21	20	19	18	
						Col	x	x	x	x	x	AP	12	11	10	9	8	7	6	5	4	3	
0011b		Reserved				Row																	
						Col																	
0100b		Reserved				Row																	
						Col																	
0101b	1GB	1Gb, x8 2Gb, x16	15	14	13	Row	x	x	17	16	29	28	27	26	25	24	23	22	21	20	19	18	
						Col	x	x	x	x	x	AP	12	11	10	9	8	7	6	5	4	3	
0110b		Reserved				Row																	
						Col																	
0111b	2GB	2Gb, x8 4Gb, x16	15	14	13	Row	x	17	16	30	29	28	27	26	25	24	23	22	21	20	19	18	
						Col	x	x	x	x	x	AP	12	11	10	9	8	7	6	5	4	3	
1000b		Reserved				Row																	
						Col																	
1001b		Reserved				Row																	
						Col																	
1010b	4GB	4Gb, x8 8Gb, x16	15	14	13	Row	17	16	31	30	29	28	27	26	25	24	23	22	21	20	19	18	
						Col	x	x	x	x	x	AP	12	11	10	9	8	7	6	5	4	3	
1011b	8GB	8Gb, x8	16	15	14	Row	17	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	
						Col	x	x	x	x	13	AP	12	11	10	9	8	7	6	5	4	3	
1100b		Reserved				Row																	
						Col																	

D18F2x84_dct[1:0] DRAM MRS

Reset: 0000_0005h. See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:24	Reserved.

23	<p>PchgPDMoSel: precharge power down mode select. Read-write. BIOS: 1. Specifies how a chip select enters and exits power down mode. This mode is enabled by D18F2x94_dct[1:0][PowerDownEn] and its behavior varies based on the setting of D18F2x94_dct[1:0][PowerDownMode] and MR0[PPD] in Table 38. See 2.9.5.8 [DRAM Device and Controller Initialization].</p> <table border="1"> <thead> <tr> <th><u>PowerDownMode</u></th> <th><u>PchgPDMoSel</u></th> <th><u>MR0[PPD]</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>0b</td> <td>0b</td> <td>Full channel slow exit (DLL off)</td> </tr> <tr> <td>0b</td> <td>0b</td> <td>1b</td> <td>Full channel fast exit (DLL on)</td> </tr> <tr> <td>0b</td> <td>1b</td> <td>xb</td> <td>Full channel dynamic fast exit/slow exit</td> </tr> <tr> <td>1b</td> <td>0b</td> <td>0b</td> <td>Reserved (Full channel slow exit)</td> </tr> <tr> <td>1b</td> <td>0b</td> <td>1b</td> <td>Reserved</td> </tr> <tr> <td>1b</td> <td>1b</td> <td>xb</td> <td>Partial channel dynamic fast exit/slow exit</td> </tr> </tbody> </table> <p>See D18F2x248_dct[1:0]_mp[1:0][Txpdll, Txp]. In dynamic fast exit/slow exit power down mode, the DCT dynamically issues MRS command(s) to the DRAM to specify the powerdown mode; the DCT specifies fast exit mode when chip selects on one of the two CKEs has recently been active; it specifies deep power down when chip selects on all CKEs have been idle.</p>	<u>PowerDownMode</u>	<u>PchgPDMoSel</u>	<u>MR0[PPD]</u>	<u>Description</u>	0b	0b	0b	Full channel slow exit (DLL off)	0b	0b	1b	Full channel fast exit (DLL on)	0b	1b	xb	Full channel dynamic fast exit/slow exit	1b	0b	0b	Reserved (Full channel slow exit)	1b	0b	1b	Reserved	1b	1b	xb	Partial channel dynamic fast exit/slow exit
<u>PowerDownMode</u>	<u>PchgPDMoSel</u>	<u>MR0[PPD]</u>	<u>Description</u>																										
0b	0b	0b	Full channel slow exit (DLL off)																										
0b	0b	1b	Full channel fast exit (DLL on)																										
0b	1b	xb	Full channel dynamic fast exit/slow exit																										
1b	0b	0b	Reserved (Full channel slow exit)																										
1b	0b	1b	Reserved																										
1b	1b	xb	Partial channel dynamic fast exit/slow exit																										
22:2	Reserved.																												
1:0	<p>BurstCtrl: burst length control. Read-write. BIOS: 01b. Specifies the number of sequential beats of DQ related to one read or write command. Requests from the processor are always 64-byte-length. Requests generated by D18F2x250_dct[1:0] are always 64-byte-length. Requests from GMC may be 32-byte or 64-byte-length. Software must ensure that GMC requests are disabled to configure the controller and drams for 8-beat burst length (e.g. during training). If this mode is changed, software must issue a mode-register set command to MR0 of the drams to place them in the same mode.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>8 beats</td> </tr> <tr> <td>01b</td> <td>dynamic 4 or 8 beats</td> </tr> <tr> <td>11b-10b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	8 beats	01b	dynamic 4 or 8 beats	11b-10b	Reserved																				
<u>Bits</u>	<u>Description</u>																												
00b	8 beats																												
01b	dynamic 4 or 8 beats																												
11b-10b	Reserved																												

D18F2x88_dct[1:0] DRAM Timing Low

See [2.9.1 \[DCT Configuration Registers\]](#).

Bits	Description														
31:30	Reserved.														
29:24	<p>MemClkDis: MEMCLK disable. Read-write. Reset: 3Fh. 1=Disable the MEMCLK. 0=Enable MEMCLK. BIOS: See 2.9.5.11. All enabled clocks should be 0; all disabled clocks should be 1.</p> <table border="1"> <thead> <tr> <th><u>Bit</u></th> <th><u>Pad</u></th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>MEMCLK[1,0]_H[0]</td> </tr> <tr> <td>[1]</td> <td>MEMCLK[1,0]_H[1]</td> </tr> <tr> <td>[2]</td> <td>MEMCLK[1,0]_H[2]</td> </tr> <tr> <td>[3]</td> <td>MEMCLK[1,0]_H[3]</td> </tr> <tr> <td>[4]</td> <td>MEMCLK[1,0]_H[4]</td> </tr> <tr> <td>[5]</td> <td>MEMCLK[1,0]_H[5]</td> </tr> </tbody> </table>	<u>Bit</u>	<u>Pad</u>	[0]	MEMCLK[1,0]_H[0]	[1]	MEMCLK[1,0]_H[1]	[2]	MEMCLK[1,0]_H[2]	[3]	MEMCLK[1,0]_H[3]	[4]	MEMCLK[1,0]_H[4]	[5]	MEMCLK[1,0]_H[5]
<u>Bit</u>	<u>Pad</u>														
[0]	MEMCLK[1,0]_H[0]														
[1]	MEMCLK[1,0]_H[1]														
[2]	MEMCLK[1,0]_H[2]														
[3]	MEMCLK[1,0]_H[3]														
[4]	MEMCLK[1,0]_H[4]														
[5]	MEMCLK[1,0]_H[5]														
23:0	Reserved.														

D18F2x8C_dct[1:0] DRAM Timing High

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers].

Bits	Description										
31:19	Reserved.										
18	DisAutoRefresh: disable automatic refresh. Read-write. BIOS: See 2.9.5.7. 1=Automatic refresh is disabled. See 2.9.5.9.1 [Write Levelization Training] and 2.9.5.9.6 [Continuous Pattern Generation].										
17:16	Tref: refresh rate. Read-write. This specifies the average time between refresh requests to all DRAM devices. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Undefined behavior.</td> </tr> <tr> <td>01b</td> <td>Reserved</td> </tr> <tr> <td>10b</td> <td>Every 7.8 us</td> </tr> <tr> <td>11b</td> <td>Every 3.9 us</td> </tr> </tbody> </table>	Bits	Description	00b	Undefined behavior.	01b	Reserved	10b	Every 7.8 us	11b	Every 3.9 us
Bits	Description										
00b	Undefined behavior.										
01b	Reserved										
10b	Every 7.8 us										
11b	Every 3.9 us										
15:0	Reserved.										

D18F2x90_dct[1:0] DRAM Configuration Low

See 2.9.1 [DCT Configuration Registers].

Bits	Description										
31:28	Reserved.										
27	DisDllShutdownSR: disable DLL shutdown in self-refresh mode. Read-write. Reset: 1. BIOS: See 2.9.5.4.2.1. 1=Disable the power saving features of shutting down DDR phy DLLs during DRAM self refresh and memory P-states. 0=Shutdown DLLs during DRAM self refresh and allow memory P-state transitions. Setting this bit does not effect the current memory P-state.										
26	Reserved.										
25	PendRefPaybackS3En: pending refresh payback S3 enable. Read-write. Reset: 0. BIOS: 1. Specifies the S3 refresh payback behavior when PendRefPayback=0. 1=Pending refreshes are paid back on S3 entry. 0=Pending refreshes are not paid back on S3 entry.										
24	StagRefEn: Stagger Refresh Enable. Read-write. Reset: 0. BIOS: 11=The DRAM controller arbitrates refreshes among chip selects based on the Tstag value . See D18F2x228_dct[1:0]. 0=DCT arbitrates among chip selects using the Trfc value. See D18F2x208_dct[1:0].										
23	ForceAutoPchg: force auto precharging. Read-write. Reset: 0. See 2.9.5.7. 1=Force auto-precharge cycles with every read or write command.										
22:21	IdleCycLowLimit: idle cycle low limit. Read-write. Reset: 0. Specifies the number of MEMCLK cycles a page is allowed to be open before it may be closed by the dynamic page close logic. This field is ignored if D18F2x90_dct[1:0][DynPageCloseEn] = 0. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>16 clocks</td> </tr> <tr> <td>01b</td> <td>32 clocks</td> </tr> <tr> <td>10b</td> <td>64 clocks</td> </tr> <tr> <td>11b</td> <td>96 clocks</td> </tr> </tbody> </table>	Bits	Description	00b	16 clocks	01b	32 clocks	10b	64 clocks	11b	96 clocks
Bits	Description										
00b	16 clocks										
01b	32 clocks										
10b	64 clocks										
11b	96 clocks										

20	DynPageCloseEn: dynamic page close enable. Read-write. Reset: 0. See 2.9.5.7. 1=The DRAM controller dynamically determines when to close open pages based on the history of that particular page and D18F2x90_dct[1:0][IdleCycLowLimit]. 0=Any open pages not auto-precharged by the DRAM controller are automatically closed after 128 clocks of inactivity.
19	Reserved.
18	PendRefPayback: pending refresh payback. Read-write. Reset: 0. BIOS: 0. 1=The DRAM controller executes all pending refresh commands before entering the self refresh state. 0=The controller enters the self refresh state regardless of the number of pending refreshes; applies to any self refresh entry if PendRefPaybackS3En=0, else any non-S3 self refresh entry.
17	EnterSelfRef: enter self refresh command. Read, write-1-only; cleared-by-hardware. Reset: 0. 1=The DRAM controller places the DRAMs into self refresh mode. The DRAM interface is tristated 1 MEMCLK after the self refresh command is issued to the DRAMs. Once entered, the DRAM interface must remain in self refresh mode for a minimum of 5 MEMCLKs. This bit is read as a 1 while the enter-self-refresh command is executing; it is read as 0 at all other times. See 2.9.5.9 [DRAM Training] for information on how to use this bit.
16	UnbuffDimm: unbuffered DIMM. Read-write or read-only, depending on the product. Reset: Product-specific. 1=The DRAM controller is connected to unbuffered DIMMs. 0=The DRAM controller is connected to registered DIMMs or LR-DIMMs.
15:12	Reserved.
11:9	Reserved.
8	Reserved.
7:2	Reserved.
1	ExitSelfRef: exit self refresh (after suspend to RAM or for DRAM training) command. Read, write-1-only; cleared-by-hardware. Reset: 0. Writing a 1 to this bit causes the DRAM controller to bring the DRAMs out of self refresh mode. It also causes the DRAM controller to issue ZQCL and MRS MR0 commands. This command should be executed by BIOS when returning from the suspend to RAM state, after the DRAM controller configuration registers are properly initialized, or when self refresh is used during DRAM training. See 2.9.5.9 [DRAM Training]. This bit is read as a 1 while the exit-self-refresh command is executing; it is read as 0 at all other times. This bit should not be set if the DCT is disabled. After using this bit during a return from suspend to RAM, BIOS must issue an additional MRS command to set MR0[PPD]=1 if Fast exit precharge powerdown mode is desired. See D18F2x94_dct[1:0][1:0][PowerDownEn] and D18F2x84_dct[1:0][PchgPDMModeSel].
0	Reserved.

D18F2x94_dct[1:0] DRAM Configuration High

See 2.9.1 [DCT Configuration Registers].

Table 132: Memory Clock Frequency Value Definition

Bits	Description
03h-00h	Reserved
04h	333 MHz. (667 MT/s)
05h	Reserved
06h	400 MHz. (800 MT/s)
09h-07h	Reserved
0Ah	533 MHz. (1066 MT/s)

Table 132: Memory Clock Frequency Value Definition

Bits	Description
0Dh-0Bh	Reserved
0Eh	667 MHz. (1333 MT/s)
11h-0Fh	Reserved
12h	800 MHz. (1600 MT/s)
15h-13h	Reserved
16h	933 MHz. (1866 MT/s)
18h-17h	Reserved
19h	1050 MHz. (2100 MT/s)
1Ah	1066 MHz. (2133 MT/s)
1Eh-1Bh	Reserved
1Fh	1200 MHz. (2400 MT/s)

Bits	Description						
31	DphyMemPsSelEn: Ddr phy MemPsSel enable. Read-write. Reset: 0h. BIOS: 1. 1=The DCT uses D18F1x10C[MemPsSel] to select the memory P-state context of a phy register when accessed by software with D18F2x98_dct[1:0][DctOffset[29:20]]==0D0h . DctOffset[6] is ignored. 0=Software accesses use DctOffset[6] to select context of those registers. If D18F2x98_dct[1:0][DctOffset[29:20]]!=0D0h then this register has no effect. See D18F2x9C_x0D04_E008_dct[1:0][PstateToAccess] .						
30:29	Reserved.						
28:24	DcqBypassMax: DRAM controller queue bypass maximum. Read-write. Reset: 0h. BIOS: 2.9.5.7 . The DRAM controller arbiter normally allows transactions to pass other transactions in order to optimize DRAM bandwidth. This field specifies the maximum number of times that the oldest memory-access request in the DRAM controller queue may be bypassed before the arbiter decision is overridden and the oldest memory-access request is serviced instead. <table border="1" data-bbox="267 1291 1448 1396"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>No bypass; the oldest request is never bypassed.</td> </tr> <tr> <td>1Fh-1h</td> <td>The oldest request may be bypassed no more than <DcqBypassMax> time.</td> </tr> </tbody> </table>	Bits	Description	0h	No bypass; the oldest request is never bypassed.	1Fh-1h	The oldest request may be bypassed no more than <DcqBypassMax> time.
Bits	Description						
0h	No bypass; the oldest request is never bypassed.						
1Fh-1h	The oldest request may be bypassed no more than <DcqBypassMax> time.						
23	ProcOdtDis: processor on-die termination disable. Read-write. Reset: 0h. 1=The processor-side on-die termination is disabled. 0=Processor-side on-die termination enabled. See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][ProcOdt] for ODT definitions. Changes to this bit must be performed prior to DRAM initialization.						

22	<p>BankSwizzleMode: bank swizzle mode. Read-write. Reset: 0. See 2.9.5.7. 1=Remaps the DRAM device bank address bits as a function of normalized physical address bits. Each of the bank address bits, as specified in Table 131 of D18F2x80_dct[1:0], are remapped as follows:</p> <ul style="list-style-type: none"> • Define X as a bank address bit (e.g., X=15 if the bank bit is specified to be address bit 15). • Define S(n) as the state of address bit n (0 or 1) and B as the remapped bank address bit. Then, <ul style="list-style-type: none"> B= S(X) ^ S(X + 3) ^ S(X + 6); for an 8-bank DRAM and D18F2x[6C:60]_dct[1:0][RankDef] = 0xb. B= S(X) ^ S(X + 4) ^ S(X + 7); for an 8-bank DRAM and D18F2x[6C:60]_dct[1:0][RankDef] = 10b. B= S(X) ^ S(X + 5) ^ S(X + 8); for an 8-bank DRAM and D18F2x[6C:60]_dct[1:0][RankDef] = 11b. <p>For example, encoding 02h of Table 131 would be remapped from Bank[2:0]={A15, A14, A13} to the following: Bank[2:0] = {A15^A18^A21, A14^A17^A20, A13^A16^A19}.</p>
21	<p>FreqChgInProg: frequency change in progress. Read-only. Reset: 0. 1=A MEMCLK frequency change is in progress. The DDR phy asserts this bit when it is in the process of locking the PLL. BIOS should not program the phy registers while this bit is set. 0=DRAM-interface commands can be sent to the phy.</p>
20	<p>SlowAccessMode: slow access mode (a.k.a. 2T mode). Read-write. Reset: 0. 1=One additional MEMCLK of setup time is provided on all DRAM address and control signals (not including CS, CKE, and ODT); i.e., these signals are driven for two MEMCLK cycles rather than one. 0=DRAM address and control signals are driven for one MEMCLK cycle. 2T mode may be needed in order to meet electrical requirements of certain DIMM speed and loading configurations. See 2.9.5.6.6 [DRAM Address Timing and Output Driver Compensation Control].</p>
19	Reserved.
18:17	Reserved.
16	<p>PowerDownMode: power down mode. Read-write. Reset: 0. BIOS: 1. Specifies how a chip select or group of chip selects enters power down mode when enabled by D18F2x94_dct[1:0][PowerDownEn]. A chip select enters power down mode when the DCT deasserts the CKE pin. The command and address signals tristate one MEMCLK after CKE deasserts. The DCT behavior varies based on the setting of D18F2x84_dct[1:0][PchgPDMoSel]. See also Table 130 [DIMM, Chip Select, and Register Mapping].</p> <p>0=Channel CKE control mode; the DRAM channel is placed in power down mode when all chip selects associated with the channel are idle; both CKE pins for the channel operate in lock step in terms of placing the channel in power down mode.</p> <p>1=Chip select CKE control mode; the chip select group controlled by a CKE pin is placed in power down mode when no transactions are pending.</p>
15	<p>PowerDownEn: power down mode enable. Read-write. Reset: 0. BIOS: 2.9.5.7. 1=Power down mode is enabled. Only precharge power down mode is supported, not active power down mode. See PowerDownMode, D18F2x84_dct[1:0][PchgPDMoSel], D18F2xA8_dct[1:0][PrtlChPDEnhEn, AggrPDEn, PDPhyPSDis], and D18F2x248_dct[1:0]_mp[1:0][PchgPDEnDelay].</p>
14	<p>DisDramInterface: disable the DRAM interface. Read-write. Reset: 0. 1=The DRAM controller is disabled and the DRAM interface is placed into a low power state. This bit must be set if there are no DIMMs connected to the DCT.</p>
13	Reserved.
12	Reserved.

11:10	<p>ZqcsInterval: ZQ calibration short interval. Read-write. Reset: 00b. BIOS: See 2.9.5.7. This field specifies the programmable interval for the controller to send out the DRAM ZQ calibration short command.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>ZQ calibration short command is disabled</td> </tr> <tr> <td>01b</td> <td>64 ms</td> </tr> <tr> <td>10b</td> <td>128 ms</td> </tr> <tr> <td>11b</td> <td>256 ms</td> </tr> </tbody> </table>	Bits	Description	00b	ZQ calibration short command is disabled	01b	64 ms	10b	128 ms	11b	256 ms
Bits	Description										
00b	ZQ calibration short command is disabled										
01b	64 ms										
10b	128 ms										
11b	256 ms										
9:8	Reserved.										
7	<p>MemClkFreqVal: memory clock frequency valid. Read-write. Reset: 0. System BIOS should set this bit after setting up D18F2x94_dct[1:0][MemClkFreq] to the proper value. This indicates to the DRAM controller that it may start driving MEMCLK at the proper frequency. After programming MemClkFreqVal=1, BIOS should poll FreqChgInProg to determine when the DRAM-interface clocks are stable. This bit should not be set if the DCT is disabled. BIOS must change each DCT's operating frequency in order. See 2.9.5.8 [DRAM Device and Controller Initialization].</p>										
6:5	Reserved.										
4:0	<p>MemClkFreq: memory clock frequency. Read-write. Reset: 000b. Specifies the frequency and rate of the DRAM interface (MEMCLK). See: Table 132 [Memory Clock Frequency Value Definition]. The rate is twice the frequency. See D18F5x84[DdrMaxRate] and D18F5x84[DdrMaxRateEnf].</p>										

D18F2x98_dct[1:0] DRAM Controller Additional Data Offset

Reset: 8000_0000h. See 2.9.1 [DCT Configuration Registers].

Each DCT includes an array of registers that are used primarily to control DRAM-interface electrical parameters. Access to these registers is accomplished as follows:

Reads:

1. Write the register number to [D18F2x98_dct\[1:0\]\[DctOffset\]](#) with [D18F2x98_dct\[1:0\]\[DctAccess-Write\]=0](#).
2. Read the register contents from [D18F2x9C_dct\[1:0\]](#).

Writes:

1. Write all 32 bits of register data to [D18F2x9C_dct\[1:0\]](#) (individual byte writes are not supported).
2. Write the register number to [D18F2x98_dct\[1:0\]\[DctOffset\]](#) with [D18F2x98_dct\[1:0\]\[DctAccess-Write\]=1](#).
 - The data will be delivered to the phy similar to a posted memory-write, and the write will complete without any further action. However, to ensure that the contents of the array register write have been delivered to the phy, software issues a subsequent configuration register read or write to any register in the northbridge. For example, reading [D18F2x98_dct\[1:0\]](#) will accomplish this.

Registers for which there is one instance per memory P-state (listed with “_mp[1:0]” appended to the register mnemonic) use [D18F1x10C\[MemPsSel\]](#), [D18F2x94_dct\[1:0\]\[DphyMemPsSelEn\]](#), and [D18F2x9C_x0D04_E008_dct\[1:0\]\[PStateToAccess\]](#) for software accesses. BIOS programs these fields appropriately to ensure consistency of registers with controller fields and DDR phy fields.

- [D18F2x9C_x0000_0\[3:0\]0\[2:1\]_dct\[1:0\]_mp\[1:0\]](#) refers to all instances of the [D18F2x9C_x0000_0\[3:0\]0\[3:1\]](#) register.
- [D18F2x9C_x0000_0\[3:0\]0\[3:1\]_dct\[1:0\]_mp\[1\]](#) refers to the register for memory P-state 1 of either or both DCTs.

- It is recommended that BIOS program context sensitive registers in batches (training/restoring all registers of a context before selecting a new context). BIOS should do the following prior to a new “batch”:
 - Program D18F2x94_dct[1:0][DphyMemPsSelEn]=1.
 - Program D18F2x9C_x0D04_E008_dct[1:0][PStateToAccess]=<target>.
 - Program D18F1x10C[MemPsSel]=<target>.

Writes to any register in this additional address space causes the FIFO pointers to be reset. Therefore, it is recommended that only BIOS write these registers. System software other than BIOS may reliably read certain memory P-state specific DDR Phy registers with DctOffset[29:20]!=0D0h in-context normally, or out-of-context using alternative direct access methods (without the requirement for programming D18F2x9C_x0D04_E008_dct[1:0][PStateToAccess]). BIOS may also use this method after memory traffic is enabled, for reading and saving register state in preparation for S3 support. The following table shows the fields that are supported and the equivalent phy direct access:

Table 133: Memory P-state Specific Indirect CSR Access by non-BIOS Software

Indirect Register	Equivalent Direct Register
D18F2x9C_x0000_0000_dct[1:0]_mp[1:0] [ProcOdt]	D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0] [ProcOdt]
D18F2x9C_x0000_0000_dct[1:0]_mp[1:0] [DqsDrvStren]	D18F2x9C_x0D0F_0[F,7:0]04_dct[1:0]_mp[1:0] [DqsDrvStren]
D18F2x9C_x0000_0000_dct[1:0]_mp[1:0] [DataDrvStren]	D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0] [DataDrvStren]
D18F2x9C_x0000_0000_dct[1:0]_mp[1:0] [ClkDrvStren]	D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0] [ClkDrvStren]
D18F2x9C_x0000_0000_dct[1:0]_mp[1:0] [AddrCmdDrvStren]	D18F2x9C_x0D0F_8[1:0]0[8,4,0]_dct[1:0]_mp[1:0] [DrvStrength]
D18F2x9C_x0000_0000_dct[1:0]_mp[1:0] [CsOdtDrvStren]	D18F2x9C_x0D0F_8[1:0]0[8,4,0]_dct[1:0]_mp[1:0] [DrvStrength]
D18F2x9C_x0000_0004_dct[1:0]_mp[1:0] [AddrCmdFineDelay]	D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0] [Delay]
D18F2x9C_x0000_0004_dct[1:0]_mp[1:0] [CsOdtFineDelay]	D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0] [Delay]
D18F2x9C_x0000_0004_dct[1:0]_mp[1:0] [CkeFineDelay]	D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0] [Delay]
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] [RxDLLWakeupTime]	D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0] [RxDLLWakeupTime]
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] [RxCPUpdPeriod]	D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0] [RxCPUpdPeriod]
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] [RxMaxDurDllNoLock]	D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0] [RxMaxDurDllNoLock]

Table 133: Memory P-state Specific Indirect CSR Access by non-BIOS Software

Indirect Register	Equivalent Direct Register
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] [TxDLLWakeupTime]	D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0] [TxDLLWakeupTime]
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] [TxCPUUpdPeriod]	D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0] [TxCPUUpdPeriod]
D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] [TxMaxDurDllNoLock]	D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0] [TxMaxDurDllNoLock]

System software other than BIOS may read any memory P-state specific DCT register with DctOffset[29:20]≠0D0h in-context normally, or out-of-context by programming D18F1x10C[MemPsSel]. The following table shows the fields that are supported. For indirect registers with both DCT and DDR Phy fields, software should mask off the DDR Phy fields as these may not be correct without programming D18F2x9C_x0D04_E008_dct[1:0][PStateToAccess].

Table 134: Memory P-state Specific DCT CSR In Indirect Space

Indirect Register
03h-00h
D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0][WrDatGrossDly] (for all bytelanes)
D18F2x9C_x0000_0004_dct[1:0]_mp[1:0][AddrCmdSetup, CsOdtSetup, CkeSetup]
D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnGrossDelay] (for all bytelanes)
D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0][WrDqsGrossDly] (for all bytelanes)

Reads or writes to any register in this additional address space collide with system self-refresh requests. Once powermanagement is enabled software should temporarily disable powermanagement prior to accessing these registers.

Bits	Description
31	Reserved.
30	DctAccessWrite: DRAM controller read/write select. RAZ; write. 0=Specifies a read access. 1=Specifies a write access.
29:0	DctOffset: DRAM controller offset. Read-write.

D18F2x9C_dct[1:0] DRAM Controller Additional Data Port

See D18F2x98_dct[1:0] for register access information. See 2.9.1 [DCT Configuration Registers].

D18F2x9C_x0000_0000_dct[1:0]_mp[1:0] DRAM Output Driver Compensation Control

See 2.9.5.6.6 [DRAM Address Timing and Output Driver Compensation Control].

Bits	Description
31	Reserved.
30:28	ProcOdt. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0][ProcOdt]. Writes to this field write all byte-lanes and groups 0 and 2. Reads to this field read byte-lane 0 group 0 ProcOdt.

27:23	Reserved.
22:20	DqsDrvStren. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]04_dct[1:0]_mp[1:0] [DqsDrvStren]. Writes to this field write all byte-lanes. Reads to this field read byte-lane 0 group 1 DqsDrvStren.
19	Reserved.
18:16	DataDrvStren. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0] [DataDrvStren]. Writes to this field write all byte-lanes and groups 0 and 2. Reads to this field read byte-lane 0 group 0 DataDrvStren.
15	Reserved.
14:12	ClkDrvStren. Read-write. DdrPhy: See D18F2x9C_x0D0F_2[F,2:0]00_dct[1:0]_mp[1:0] [Clk-DrvStren]. Writes to this field write all clock chips. Reads to this field read clock chip 0 ClkDrvStren.
11	Reserved.
10:8	AddrCmdDrvStren. Read-write. DdrPhy: See D18F2x9C_x0D0F_8[1:0]0[8,4,0]_dct[1:0]_mp[1:0] [DrvStrength] and D18F2x9C_x0D0F_C0[10,0C,08,04]_dct[1:0]_mp[1:0] [DrvStrength]. Writes to this field write all groups of Addr/Cmd 1 chip and all groups of Address2 chip. Reads to this field read DrvStrength of group 0 of Addr/Cmd 1 chip.
7	Reserved.
6:4	CsOdtDrvStren. Read-write. DdrPhy: See D18F2x9C_x0D0F_8[1:0]0[8,4,0]_dct[1:0]_mp[1:0] [DrvStrength]. Writes to this field write all groups of Addr/Cmd 0 chip. Reads to this field read DrvStrength of group 0 of Addr/Cmd 0 chip.
3	Reserved.
2:0	CkeDrvStren. Read-write. DdrPhy: See D18F2x9C_x0D0F_C000_dct[1:0]_mp[1:0] [CkeDrvStren].

D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0] DRAM Write Data Timing

BIOS: See 2.9.5.9.4 [DQS Position Training].

Table 135: Index Mapping for D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:0]	Function ¹
0000_0001h	DIMM/CS 0 Bytes 3-0
0000_0002h	DIMM/CS 0 Bytes 7-4
0000_0101h	DIMM/CS 1 Bytes 3-0
0000_0102h	DIMM/CS 1 Bytes 7-4
0000_0201h	CS 2 Bytes 3-0
0000_0202h	CS 2 Bytes 7-4
0000_0301h	CS 3 Bytes 3-0
0000_0302h	CS 3 Bytes 7-4

1. If [D18F2xA8_dct\[1:0\]](#)[PerRankTimingEn]=1 then the function is CS. Otherwise the function is DIMM.

Table 136: Byte Lane Mapping for D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0]

Register	Bits			
	31:24	23:16	15:8	7:0
D18F2x9C_x0000_0[3:0]01	Byte3	Byte2	Byte1	Byte0
D18F2x9C_x0000_0[3:0]02	Byte7	Byte6	Byte5	Byte4

These registers control the timing of write DQ with respect to MEMCLK and allow transmit DQS to be centered in the data eye. The delay starts 1 UI before the rising edge of MEMCLK corresponding to the CAS-write-latency. See 2.9.5.9 [DRAM Training]. WrDatGrossDly must be programmed for a given DIMM and lane such that WrDatDly - WrDqsDly \leq 0.5 MEMCLKs.

Bits	Description										
31:29	WrDatGrossDly3: write data gross delay. See: D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0][WrDatGrossDly] .										
28:24	WrDatFineDly3: write data fine delay. See: D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0][WrDatFineDly] .										
23:21	WrDatGrossDly2: write data gross delay. See: D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0][WrDatGrossDly] .										
20:16	WrDatFineDly2: write data fine delay. See: D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0][WrDatFineDly] .										
15:13	WrDatGrossDly1: write data gross delay. See: D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0][WrDatGrossDly] .										
12:8	WrDatFineDly1: write data fine delay. See: D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0][WrDatFineDly] .										
7:5	WrDatGrossDly: write data gross delay. Read-write. Reset: 0. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>No delay</td> </tr> <tr> <td>001b</td> <td>0.5 MEMCLK delay</td> </tr> <tr> <td>110b-010b</td> <td><WrDatGrossDly>/2 MEMCLK delay</td> </tr> <tr> <td>111b</td> <td>3.5 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	000b	No delay	001b	0.5 MEMCLK delay	110b-010b	<WrDatGrossDly>/2 MEMCLK delay	111b	3.5 MEMCLK delay
Bits	Description										
000b	No delay										
001b	0.5 MEMCLK delay										
110b-010b	<WrDatGrossDly>/2 MEMCLK delay										
111b	3.5 MEMCLK delay										
4:0	WrDatFineDly: write data fine delay. Read-write. Cold reset: 0. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0/64 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><WrDatFineDly>/64 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	00h	0/64 MEMCLK delay	1Eh-01h	<WrDatFineDly>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay		
Bits	Description										
00h	0/64 MEMCLK delay										
1Eh-01h	<WrDatFineDly>/64 MEMCLK delay										
1Fh	31/64 MEMCLK delay										

D18F2x9C_x0000_0004_dct[1:0]_mp[1:0] DRAM Address/Command Timing Control

BIOS: 2.9.5.6.6. This register controls the timing of the address, command, chip select, ODT and clock enable pins with respect to MEMCLK as shown in Figure 11. See 2.9.5.4.2 [DRAM Channel Frequency Change] and 2.9.5.4.3 [Phy Fence Programming].

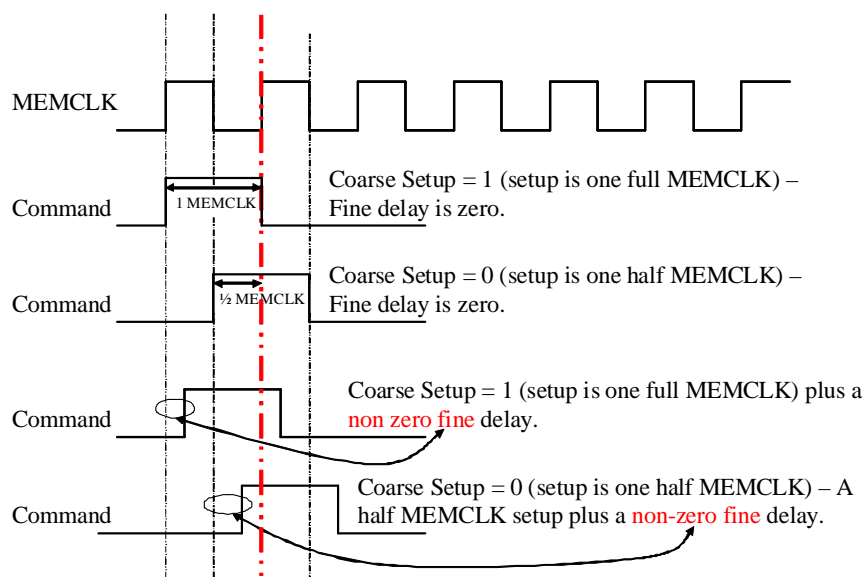


Figure 11: Address/Command Timing at the Processor Pins

2T timing is controlled by `D18F2x94_dct[1:0][SlowAccessMode]`. If a setup time (course delay) field is changed and `D18F2x94_dct[1:0][MemClkFreqVal]=1`, then software must toggle `MemClkFreqVal` for the delay to take effect.

Bits	Description
31:22	Reserved.
21	AddrCmdSetup: address/command setup time. Read-write. Reset: 0. DCT: Selects the default setup time for the address and command pins versus MEMCLK. 0=1/2 MEMCLK (1 1/2 MEMCLK for 2T timing). 1=1 MEMCLK (2 MEMCLKs for 2T timing).
20:16	AddrCmdFineDelay. Read-write. Cold Reset: 0. DdrPhy: See D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Delay] and D18F2x9C_x0D0F_C020_dct[1:0]_mp[1:0][Delay] . Writes to this field write D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Delay] of the Cmd/Addr 1 chip and to D18F2x9C_x0D0F_C020_dct[1:0]_mp[1:0][Delay] . Reads to this field read D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Delay] from Cmd/Addr 1 chip.
15:14	Reserved.
13	CsOdtSetup: CS/ODT setup time. Read-write. Reset: 0. DCT: Selects the default setup time for the CS and ODT pins versus MEMCLK. 0=1/2 MEMCLK. 1=1 MEMCLK.
12:8	CsOdtFineDelay. Read-write. Cold Reset: 0. DdrPhy: See D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Delay] . Writes to this field write Delay of the Cmd/Addr 0 chip. Reads to this field read Delay from Cmd/Addr 0 chip.
7:6	Reserved.
5	CkeSetup: CKE setup time. Read-write. Reset: 0. DCT: Selects the default setup time for the CKE pins versus MEMCLK. 0=1/2 MEMCLK. 1=1 MEMCLK.
4:0	CkeFineDelay: CKE fine delay. Read-write. Cold reset: 00h. DdrPhy: See D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0][Delay] . Writes to this field write to D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0][Delay] . Reads to this field read from D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0][Delay] .

D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0] DRAM Read DQS Timing**Table 137: Index Mapping for D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0]**

D18F2x98_dct[1:0][31:0]	Function ¹
0000_0005h	DIMM/CS 0 Bytes 3-0
0000_0006h	DIMM/CS 0 Bytes 7-4
0000_0105h	DIMM/CS 1 Bytes 3-0
0000_0106h	DIMM/CS 1 Bytes 7-4
0000_0205h	Reserved /CS 2 Bytes 3-0
0000_0206h	Reserved /CS 2 Bytes 7-4
0000_0305h	Reserved /CS 3 Bytes 3-0
0000_0306h	Reserved /CS 3 Bytes 7-4

1. If D18F2xA8_dct[1:0][PerRankTimingEn]=1 then the function is CS. Otherwise the function is DIMM.

Table 138: Byte Lane Mapping for D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0]

Register	Bits			
	29:25	21:17	13:9	5:1
D18F2x9C_x0000_0[3:0]05	Byte3	Byte2	Byte1	Byte0
D18F2x9C_x0000_0[3:0]06	Byte7	Byte6	Byte5	Byte4

These registers control the timing of read (input) DQS signals with respect to DQ. See 2.9.5.9 [DRAM Training]. The actual delay applied to the DQS input signal before sampling data includes an internal part dependent delay plus the nominal register delay specified here. The part dependent (insertion) delay is large in proportion to individual step delay controlled by this register. If the actual delay exceeds 1UI and read bubbles exist in the read data stream then incorrect operation of received DQS may occur. BIOS training allows this scenario temporarily by ensuring no bubbles in read streams exist. The actual delay of the optimally trained RdDqsTime value does not exceed 1UI.

Bits	Description
31:30	Reserved.
29:25	RdDqsTime3: read DQS timing control. See: D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0][RdDqsTime] .
24:22	Reserved.
21:17	RdDqsTime2: read DQS timing control. See: D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0][RdDqsTime] .
16:14	Reserved.
13:9	RdDqsTime1: read DQS timing control. See: D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0][RdDqsTime] .
8:6	Reserved.

5:1	RdDqsTime: read DQS timing control. Read-write. Cold reset: 0Fh. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0/64 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><RdDqsTime>/64 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00h	0/64 MEMCLK delay	1Eh-01h	<RdDqsTime>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay
<u>Bits</u>	<u>Description</u>								
00h	0/64 MEMCLK delay								
1Eh-01h	<RdDqsTime>/64 MEMCLK delay								
1Fh	31/64 MEMCLK delay								
0	Reserved.								

D18F2x9C_x0000_0008_dct[1:0]_mp[1:0] DRAM Phy Control

Cold reset: 0208_0000h. See [2.9.5.9 \[DRAM Training\]](#).

Bits	Description										
31	Reserved.										
30	DisAutoComp: disable automatic compensation. Read-write. 1=Disable the compensation control state machine. 0=The phy automatic compensation engine is enabled.										
29	DisablePredriverCal: disable predriver calibration. Read-write. BIOS: See 2.9.5.4.4 . 1=Disable the update of predriver codes to all pads. Setting this bit in DCT0 disables calibration engine for DCT0 and DCT1. Setting this bit in DCT1 has no effect.										
28:14	Reserved.										
13	DqsRcvTrEn: DQS receiver training enable. Read-write. 1=Initiate hardware assisted read DQS receiver training. 0=Stop read DQS receiver training. This allows BIOS to reliably read the DQS receiver training data.										
12	WrLvOdtEn: write levelization ODT enabled. Read-write. 1=ODT specified by WrLvOdt is enabled during write levelization training. 0=ODT is disabled during write levelization training.										
11:8	WrLvOdt: write levelization ODT. Read-write. Specifies the state of the ODT pins when WrLvOdtEn is set. 1=ODT is enabled. 0=ODT is disabled. See 2.9.5.6.5 [DRAM ODT Control] . Tri-state enable for ODT is turned off by the phy when WrLvOdtEn=1.										
7:6	FenceTrSel: fence train select. Read-write. Specifies the flop to be used for phy based fence training. See PhyFenceTrEn. This field is shared by D18F2x9C_x0000_0008_dct[1:0]_mp0 and D18F2x9C_x0000_0008_dct[1:0]_mp1. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>PRE flop</td> </tr> <tr> <td>01b</td> <td>RxDll flop</td> </tr> <tr> <td>10b</td> <td>TxDll flop</td> </tr> <tr> <td>11b</td> <td>TxPad flop</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	PRE flop	01b	RxDll flop	10b	TxDll flop	11b	TxPad flop
<u>Bits</u>	<u>Description</u>										
00b	PRE flop										
01b	RxDll flop										
10b	TxDll flop										
11b	TxPad flop										

5:4	<p>TrChipSel: training DIMM select. Read-write. Specifies a timing control for a corresponding rank which (see D18F2xA8_dct[1:0][PerRankTimingEn]) is to be trained.</p> <p>IF (D18F2xA8_dct[1:0][PerRankTimingEn]==0) THEN</p> <table border="0"> <tr> <td><u>Bits</u></td> <td><u>Description</u></td> </tr> <tr> <td>00b</td> <td>chip select 0 and 1</td> </tr> <tr> <td>01b</td> <td>chip select 2 and 3</td> </tr> <tr> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </table> <p>ELSE</p> <table border="0"> <tr> <td><u>Bits</u></td> <td><u>Description</u></td> </tr> <tr> <td>00b</td> <td>chip select 0</td> </tr> <tr> <td>01b</td> <td>chip select 1</td> </tr> <tr> <td>10b</td> <td>chip select 2</td> </tr> <tr> <td>11b</td> <td>chip select 3</td> </tr> </table> <p>ENDIF.</p>	<u>Bits</u>	<u>Description</u>	00b	chip select 0 and 1	01b	chip select 2 and 3	10b	Reserved	11b	Reserved	<u>Bits</u>	<u>Description</u>	00b	chip select 0	01b	chip select 1	10b	chip select 2	11b	chip select 3
<u>Bits</u>	<u>Description</u>																				
00b	chip select 0 and 1																				
01b	chip select 2 and 3																				
10b	Reserved																				
11b	Reserved																				
<u>Bits</u>	<u>Description</u>																				
00b	chip select 0																				
01b	chip select 1																				
10b	chip select 2																				
11b	chip select 3																				
3	PhyFenceTrEn: phy fence training enable. Read-write. 1=Initiate phy based fence training. 0=Stop the phy based fence training engine.																				
2	TrNibbleSel: training nibble select. Read-write. BIOS: 0. Specifies nibbles of each DIMM data byte trained during write levelization training. 0=Lower nibbles. 1=Upper nibbles.																				
1	Reserved.																				
0	WrtLvTrEn: write levelization training enable. Read-write. 1=Initiate write levelization (tDQSS margining) training. 0=The phy stops driving DQS and exits write levelization training.																				

D18F2x9C_x0000_000B_dct[1:0] DRAM Phy Status Register

Bits	Description
31	DynModeChange: dynamic mode change. RAZ; write. Reset: 0. 1=Phy enters the state specified by PhySelfRefreshMode .
30	PhyPSReq: phy pstate request. RAZ; write. Reset: 0. 1=Phy enters the memory P-state specified by PhyPS .
29:27	Reserved.
26	PhyPS: phy pstate. RAZ; write. Reset: 0. 1=M1. 0=M0. See PhyPSReq .
25:24	Reserved
23	PhySelfRefreshMode: phy self refresh mode. RAZ; write. Reset: 0. 1=Enter self refresh mode. 0=Exit self refresh mode. See DynModeChange .
22:0	Reserved

D18F2x9C_x0000_000C_dct[1:0] DRAM Phy Miscellaneous

This register provides access to the DDR phy to control signal tri-state functionality. See [Table 130](#) for processor pin map. Based on the system configuration, BIOS may tri-state signals with associated chip selects that are unpopulated in an effort to conserve power. The recommendations for tri-state of ODT pins are as follows:

- BIOS tri-states ODT pins to unpopulated DIMM slots.
- BIOS tri-states ODT pins that are unused on slots with unbuffered DIMMs.
- BIOS does not tri-state any ODT pins to slots with registered DIMMs .
- For single rank registered DIMMs with address parity capability, BIOS must not tri-state the chip select pin

corresponding to the second chip select of the DIMM.

Bits	Description																		
31	Reserved.																		
30:26	FenceThresholdTxDll: phy fence threshold transmit DLL. Read-write. Cold reset: 13h. BIOS: See 2.9.5.4.3 . This field specifies the fence delay threshold value used for DQS receiver valid. See FenceThresholdTxPad.																		
25:21	FenceThresholdRxDll: phy fence threshold DQS receiver enable. Read-write. Cold reset: 13h. BIOS: See 2.9.5.4.3 . This field specifies the fence delay threshold value used for DQS receiver enable. See FenceThresholdTxPad.																		
20:16	FenceThresholdTxPad: phy fence threshold transmit pad. Read-write. Cold reset: 13h. BIOS: See 2.9.5.4.3 . This field specifies the fence delay threshold value used for write data, write DQS, Addr/Cmd, CS, ODT, and CKE. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0/64 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><FenceThresholdTxPad>/64 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	00h	0/64 MEMCLK delay	1Eh-01h	<FenceThresholdTxPad>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay										
Bits	Description																		
00h	0/64 MEMCLK delay																		
1Eh-01h	<FenceThresholdTxPad>/64 MEMCLK delay																		
1Fh	31/64 MEMCLK delay																		
15:12	CKETri: CKE tri-state. Read-write. Cold reset: 0h. 0=The CKE signal is not tri-stated. 1=Tri-state unconnected CKE signal from the processor. See 2.9.2 [DDR Pad to Processor Pin Mapping] . <table border="1"> <thead> <tr> <th>Bit</th> <th>Pad</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>MEMCKE[1:0][0]</td> </tr> <tr> <td>[1]</td> <td>MEMCKE[1:0][1]</td> </tr> <tr> <td>[2]</td> <td>MEMCKE[1:0][2]</td> </tr> <tr> <td>[3]</td> <td>MEMCKE[1:0][3]</td> </tr> </tbody> </table>	Bit	Pad	[0]	MEMCKE[1:0][0]	[1]	MEMCKE[1:0][1]	[2]	MEMCKE[1:0][2]	[3]	MEMCKE[1:0][3]								
Bit	Pad																		
[0]	MEMCKE[1:0][0]																		
[1]	MEMCKE[1:0][1]																		
[2]	MEMCKE[1:0][2]																		
[3]	MEMCKE[1:0][3]																		
11:8	ODTTri: ODT tri-state. Read-write. Cold reset: 0h. 0=The ODT signals are not tri-stated unless directed to by the DCT. 1=Tri-state unconnected ODT signals from the processor. See 2.9.2 [DDR Pad to Processor Pin Mapping] . <table border="1"> <thead> <tr> <th>Bit</th> <th>Pad</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>MEMODT[1:0][0]</td> </tr> <tr> <td>[1]</td> <td>MEMODT[1:0][1]</td> </tr> <tr> <td>[2]</td> <td>MEMODT[1:0][2]</td> </tr> <tr> <td>[3]</td> <td>MEMODT[1:0][3]</td> </tr> </tbody> </table>	Bit	Pad	[0]	MEMODT[1:0][0]	[1]	MEMODT[1:0][1]	[2]	MEMODT[1:0][2]	[3]	MEMODT[1:0][3]								
Bit	Pad																		
[0]	MEMODT[1:0][0]																		
[1]	MEMODT[1:0][1]																		
[2]	MEMODT[1:0][2]																		
[3]	MEMODT[1:0][3]																		
7:0	ChipSelTri: chip select tri-state. Read-write. Cold reset: 00h. 0=The chip select signals are not tri-stated unless directed to by the DCT. 1=Tri-state unpopulated chip selects when motherboard termination is available. See 2.9.2 [DDR Pad to Processor Pin Mapping] . <table border="1"> <thead> <tr> <th>Bit</th> <th>Pad</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>MEMCS[1:0]_L[0]</td> </tr> <tr> <td>[1]</td> <td>MEMCS[1:0]_L[1]</td> </tr> <tr> <td>[2]</td> <td>MEMCS[1:0]_L[2]</td> </tr> <tr> <td>[3]</td> <td>MEMCS[1:0]_L[3]</td> </tr> <tr> <td>[4]</td> <td>MEMCS[1:0]_L[4]</td> </tr> <tr> <td>[5]</td> <td>MEMCS[1:0]_L[5]</td> </tr> <tr> <td>[6]</td> <td>MEMCS[1:0]_L[6]</td> </tr> <tr> <td>[7]</td> <td>MEMCS[1:0]_L[7]</td> </tr> </tbody> </table>	Bit	Pad	[0]	MEMCS[1:0]_L[0]	[1]	MEMCS[1:0]_L[1]	[2]	MEMCS[1:0]_L[2]	[3]	MEMCS[1:0]_L[3]	[4]	MEMCS[1:0]_L[4]	[5]	MEMCS[1:0]_L[5]	[6]	MEMCS[1:0]_L[6]	[7]	MEMCS[1:0]_L[7]
Bit	Pad																		
[0]	MEMCS[1:0]_L[0]																		
[1]	MEMCS[1:0]_L[1]																		
[2]	MEMCS[1:0]_L[2]																		
[3]	MEMCS[1:0]_L[3]																		
[4]	MEMCS[1:0]_L[4]																		
[5]	MEMCS[1:0]_L[5]																		
[6]	MEMCS[1:0]_L[6]																		
[7]	MEMCS[1:0]_L[7]																		

D18F2x9C_x0000_000D_dct[1:0]_mp[1:0] DRAM Phy DLL Control

Cold Reset: 0000_0000h. This register defines programmable options for the phy's DLLs for power savings. There are two identical sets of configuration registers: one for the transmit DLLs (those running off of the phy's

internal PCLK which is running at rate of 2*MEMCLK) and receive DLLs (those running off of the DQS from the DIMMs). These values are programmed by BIOS based on programmed DDR frequency. This register must be programmed before DRAM device initialization.

Bits	Description
31:26	Reserved.
25:24	RxDLLWakeupTime: receive DLL wakeup time. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0][RxDLLWakeupTime] .
23	Reserved.
22:20	RxCPUpdPeriod: receive charge pump period. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0][RxCPUpdPeriod] .
19:16	RxMaxDurDllNoLock: receive maximum duration DLL no lock. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0][RxMaxDurDllNoLock] .
15:10	Reserved.
9:8	TxDLLWakeupTime: transmit DLL wakeup time. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0][TxDLLWakeupTime] .
7	Reserved.
6:4	TxCPUpdPeriod: transmit charge pump DLL wakeup time. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0][TxCPUpdPeriod] .
3:0	TxMaxDurDllNoLock: transmit maximum duration DLL no lock. Read-write. DdrPhy: See D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0][TxMaxDurDllNoLock] .

D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0] DRAM DQS Receiver Enable Timing

Table 139: Index Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:0]	Function ¹
0000_0010h	DIMM/CS 0 Bytes 1-0
0000_0011h	DIMM/CS 0 Bytes 3-2
0000_0013h	DIMM/CS 1 Bytes 1-0
0000_0014h	DIMM/CS 1 Bytes 3-2
0000_0016h	CS 2 Bytes 1-0
0000_0017h	CS 2 Bytes 3-2
0000_0019h	CS 3 Bytes 1-0
0000_001Ah	CS 3 Bytes 3-2
0000_0020h	DIMM/CS 0 Bytes 5-4
0000_0021h	DIMM/CS 0 Bytes 7-6
0000_0023h	DIMM/CS 1 Bytes 5-4
0000_0024h	DIMM/CS 1 Bytes 7-6
0000_0026h	CS 2 Bytes 5-4
0000_0027h	CS 2 Bytes 7-6
0000_0029h	CS 3 Bytes 5-4
0000_002Ah	CS 3 Bytes 7-6

1. If [D18F2xA8_dct\[1:0\]\[PerRankTimingEn\]=1](#) then the function is CS. Otherwise the function is DIMM.

Table 140: Byte Lane Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0]

Register	Bits	
	25:16	9:0
D18F2x9C_x0000_001[9,6,3,0]	Byte1	Byte0
D18F2x9C_x0000_001[A,7,4,1]	Byte3	Byte2
D18F2x9C_x0000_002[9,6,3,0]	Byte5	Byte4
D18F2x9C_x0000_002[A,7,4,1]	Byte7	Byte6

Each of these registers control the timing of the receiver enable from the start of the read preamble with respect to MEMCLK. See 2.9.5.9 [DRAM Training]. These delay registers must be programmed such that across all DIMMs and lanes $\text{MAX}(\text{DqsRcvEnDelay}) - \text{MIN}(\text{DqsRcvEnDelay}) \leq 7 \text{ UI}$.

Bits	Description										
31:26	Reserved.										
25:21	DqsRcvEnGrossDelay1: DQS receiver enable gross delay. See: D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnGrossDelay] .										
20:16	DqsRcvEnFineDelay1: DQS receiver enable fine delay. See: D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0][DqsRcvEnFineDelay] .										
15:10	Reserved.										
9:5	DqsRcvEnGrossDelay: DQS receiver enable gross delay. Read-write. Reset: 01h. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>No delay</td> </tr> <tr> <td>01h</td> <td>0 MEMCLK delay</td> </tr> <tr> <td>1Eh-02h</td> <td><DqsRcvEnGrossDelay>/2 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>15.5 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	00h	No delay	01h	0 MEMCLK delay	1Eh-02h	<DqsRcvEnGrossDelay>/2 MEMCLK delay	1Fh	15.5 MEMCLK delay
Bits	Description										
00h	No delay										
01h	0 MEMCLK delay										
1Eh-02h	<DqsRcvEnGrossDelay>/2 MEMCLK delay										
1Fh	15.5 MEMCLK delay										
4:0	DqsRcvEnFineDelay: DQS receiver enable fine delay. Read-write. Cold reset: 00h. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0/64 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><DqsRcvEnFineDelay>/64 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	00h	0/64 MEMCLK delay	1Eh-01h	<DqsRcvEnFineDelay>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay		
Bits	Description										
00h	0/64 MEMCLK delay										
1Eh-01h	<DqsRcvEnFineDelay>/64 MEMCLK delay										
1Fh	31/64 MEMCLK delay										

D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0] DRAM DQS Write Timing

BIOS: See 2.9.5.9 [DRAM Training]. Each of these registers control the DQS timing delay for write commands relative to MEMCLK. The delay starts at the rising edge of MEMCLK corresponding to the CAS-writelatency. Each control includes a gross timing field and a fine timing field, the sum of which is the total delay.

Table 141: Index Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:0]	Function ¹
0000_0030h	DIMM/CS 0 Bytes 1-0
0000_0031h	DIMM/CS 0 Bytes 3-2
0000_0033h	DIMM/CS 1 Bytes 1-0
0000_0034h	DIMM/CS 1 Bytes 3-2
0000_0036h	CS 2 Bytes 1-0
0000_0037h	CS 2 Bytes 3-2

Table 141: Index Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0]

0000_0039h	CS 3 Bytes 1-0
0000_003Ah	CS 3 Bytes 3-2
0000_0040h	DIMM/CS 0 Bytes 5-4
0000_0041h	DIMM/CS 0 Bytes 7-6
0000_0043h	DIMM/CS 1 Bytes 5-4
0000_0044h	DIMM/CS 1 Bytes 7-6
0000_0046h	CS 2 Bytes 5-4
0000_0047h	CS 2 Bytes 7-6
0000_0049h	CS 3 Bytes 5-4
0000_004Ah	CS 3 Bytes 7-6
1. If D18F2xA8_dct[1:0][PerRankTimingEn]=1 then the function is CS. Otherwise the function is DIMM.	

Table 142: Byte Lane Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0]

Register	Bits	
	23:16	7:0
D18F2x9C_x0000_003[9,6,3,0]	Byte1	Byte0
D18F2x9C_x0000_003[A,7,4,1]	Byte3	Byte2
D18F2x9C_x0000_004[9,6,3,0]	Byte5	Byte4
D18F2x9C_x0000_004[A,7,4,1]	Byte7	Byte6

Each of these registers control the DQS timing delay for write commands relative to MEMCLK. See 2.9.5.9 [DRAM Training].

Bits	Description
31:29	Reserved.
28:24	Reserved.
23:21	WrDqsGrossDly1: DQS write gross delay. See: D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0][WrDqsGrossDly] .
20:16	WrDqsFineDly1: DQS write fine delay. See: D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0][WrDqsFineDly] .
15:13	Reserved.
12:8	Reserved.

7:5	WrDqsGrossDly: DQS write gross delay. Read-write. Reset: 0.
	<u>Bits</u> <u>Description</u>
	000b No delay
	001b 0.5 MEMCLK delay
	110b-010b <WrDqsGrossDly>/2 MEMCLK delay
	111b 3.5 MEMCLK delay
4:0	WrDqsFineDly: DQS write fine delay. Read-write. Cold reset: 0.
	<u>Bits</u> <u>Description</u>
	00h 0/64 MEMCLK delay
	1Eh-01h <WrDqsFineDly>/64 MEMCLK delay
	1Fh 31/64 MEMCLK delay

D18F2x9C_x0000_00[51:50]_dct[1:0] DRAM Phase Recovery Control

These registers are used by BIOS for hardware assisted DRAM training. Writes to these registers seed the phase recovery engine prior to training. Reads from the registers indicate how much the phase recovery engine has advanced to align the MEMCLK and DQS edges and is under hardware control. See [2.9.5.9 \[DRAM Training\]](#).

Table 143: Register Mapping for D18F2x9C_x0000_00[51:50]_dct[1:0]

Register	Function
D18F2x9C_x0000_0050_dct[1:0]	Bytes 3-0
D18F2x9C_x0000_0051_dct[1:0]	Bytes 7-4

Table 144: Byte Lane Mapping for D18F2x9C_x0000_00[51:50]_dct[1:0]

Register	Bits			
	31:24	23:16	15:8	7:0
D18F2x9C_x0000_0050	Byte3	Byte2	Byte1	Byte0
D18F2x9C_x0000_0051	Byte7	Byte6	Byte5	Byte4

Bits	Description
31	Reserved.
30:29	PhRecGrossDly: phase recovery gross delay. See: PhRecGrossDly.
28:24	PhRecFineDly: phase recovery fine delay. See: PhRecFineDly.
23	Reserved.
22:21	PhRecGrossDly: phase recovery gross delay. See: PhRecGrossDly.
20:16	PhRecFineDly: phase recovery fine delay. See: PhRecFineDly.
15	Reserved.
14:13	PhRecGrossDly: phase recovery gross delay. See: PhRecGrossDly.
12:8	PhRecFineDly: phase recovery fine delay. See: PhRecFineDly.
7	Reserved.

6:5	PhRecGrossDly: phase recovery gross delay. Read-write. Reset: X. Gross timing indicates the number of half-MEMCLK periods that the phase recovery engine advanced while aligning edges.										
	<table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No delay</td> </tr> <tr> <td>01b</td> <td>0.5 MEMCLK delay</td> </tr> <tr> <td>10b</td> <td>1.0 MEMCLK delay</td> </tr> <tr> <td>11b</td> <td>1.5 MEMCLK delay</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	No delay	01b	0.5 MEMCLK delay	10b	1.0 MEMCLK delay	11b	1.5 MEMCLK delay
<u>Bits</u>	<u>Description</u>										
00b	No delay										
01b	0.5 MEMCLK delay										
10b	1.0 MEMCLK delay										
11b	1.5 MEMCLK delay										
4:0	PhRecFineDly: phase recovery fine delay. Read-write. Reset: X.										
	<table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0/64 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><PhRecFineDly>/64 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00h	0/64 MEMCLK delay	1Eh-01h	<PhRecFineDly>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay		
<u>Bits</u>	<u>Description</u>										
00h	0/64 MEMCLK delay										
1Eh-01h	<PhRecFineDly>/64 MEMCLK delay										
1Fh	31/64 MEMCLK delay										

D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0] Data Byte Pad Configuration

Cold reset: xxxx_xxxxh. BIOS: See D18F2x9C_x0000_0000_dct[1:0]_mp[1:0]. See 2.9.2.1 for group to pin mapping.

Table 145: Index Mapping for D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0]

Address Bits	Valid Values	Name
D18F2x98_dct[1:0][3:2]	2h,0h	Pad Group
D18F2x98_dct[1:0][6]	[1:0]	Memory Pstate
D18F2x98_dct[1:0][11:8]	[7:0]	Bytelane
D18F2x98_dct[1:0][11:8]	Fh	Bytelane [7:0]

Bits	Description												
31:16	Reserved.												
15:7	Reserved.												
6:4	ProcOdt: processor on-die termination. Read-write. Specifies the resistance of the on-die termination resistors. This field is valid only when D18F2x9C_x0D0F_0[F,7:0]04_dct[1:0]_mp[1:0][POdt-Off]=0.												
	<table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>240 ohms +/- 20%</td> </tr> <tr> <td>001b</td> <td>120 ohms +/- 20%</td> </tr> <tr> <td>010b</td> <td>80 ohms +/- 20%</td> </tr> <tr> <td>011b</td> <td>60 ohms +/- 20%</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	000b	240 ohms +/- 20%	001b	120 ohms +/- 20%	010b	80 ohms +/- 20%	011b	60 ohms +/- 20%	111b-100b	Reserved
<u>Bits</u>	<u>Description</u>												
000b	240 ohms +/- 20%												
001b	120 ohms +/- 20%												
010b	80 ohms +/- 20%												
011b	60 ohms +/- 20%												
111b-100b	Reserved												

3	Reserved.												
2:0	DataDrvStren: data drive strength. Read-write. This field specifies the drive strength of the DRAM data pins. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0.75x</td> </tr> <tr> <td>001b</td> <td>1.0x</td> </tr> <tr> <td>010b</td> <td>1.25x</td> </tr> <tr> <td>011b</td> <td>1.5x</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	0.75x	001b	1.0x	010b	1.25x	011b	1.5x	111b-100b	Reserved
Bits	Description												
000b	0.75x												
001b	1.0x												
010b	1.25x												
011b	1.5x												
111b-100b	Reserved												

D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0] Data Byte Transmit PreDriver Calibration

Cold reset: xxxx_xxxxh. BIOS: See 2.9.5.4.4.

Table 146: Index Mapping for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]							
	0702h	0602h	0502h	0402h	0302h	0202h	0102h	0002h
0D0Fh	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 147: Broadcast Mapping for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F02h
0D0Fh	D18F2x9C_x0D0F_0[7:0]02

Table 148: Valid Values for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP]

Bits	Description
0h	Slew Rate 0 (slowest)
8h-1h	Reserved
9h	Slew Rate 1
11h-Ah	Reserved
12h	Slew Rate 2
1Ah-13h	Reserved
1Bh	Slew Rate 3
23h-1Ch	Reserved
24h	Slew Rate 4
2Ch-25h	Reserved
2Dh	Slew Rate 5
35h-2Eh	Reserved
36h	Slew Rate 6
3Eh-37h	Reserved
3Fh	Slew Rate 7 (fastest)

Bits	Description
31:16	Reserved.
15	ValidTxAndPre: predriver calibration code valid. Read-write; cleared-by-hardware. 1=Predriver calibration codes are copied from this register and D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0] into the associated transmit pad. Hardware clears this field after the copy is complete.
14:12	Reserved.
11:6	TxPreP: PMOS predriver calibration code. Read-write. Specifies the rising edge slew rate of the transmit pad. See: Table 148 [Valid Values for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP]] . After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.
5:0	TxPreN: NMOS predriver calibration code. Read-write. Specifies the falling edge slew rate of the transmit pad. See: Table 148 [Valid Values for D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP]] . After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.

D18F2x9C_x0D0F_0[F,7:0]04_dct[1:0]_mp[1:0] Data Byte Pad Configuration

Cold reset: 0000_0033h. BIOS: See [D18F2x9C_x0000_0000_dct\[1:0\]_mp\[1:0\]](#). See 2.9.2.1 for group to pin mapping.

Table 149: [Index Mapping for D18F2x9C_x0D0F_0\[F,7:0\]04_dct\[1:0\]_mp\[1:0\]](#)

Address Bits	Valid Values	Name
D18F2x98_dct[1:0][3:2]	1h	Dbyte Pad Group
D18F2x98_dct[1:0][6]	[1:0]	Memory Pstate
D18F2x98_dct[1:0][11:8]	[7:0]	Bytelane
D18F2x98_dct[1:0][11:8]	Fh	Bytelane [7:0]

Bits	Description												
31:13	Reserved.												
12	POdtOff: processor ODT off. Read-write. BIOS: See 2.9.5.6.6.1=Phy disables receiver pad termination. 0=Phy enables receiver pad termination. Hardware programs this field for the current M-state (D18F2x9C_x0D0F_E006_dct[1:0][PhyPS]) with the value in D18F2x94_dct[1:0][ProcOdtDis] when the memory frequency is updated. See 2.9.5.4.2 . BIOS must reprogram this field after each frequency change if the target value differs from D18F2x94_dct[1:0][ProcOdtDis] .												
11:7	Reserved.												
6:4	ProcOdt: processor on-die termination. Read-write. Specifies the resistance of the on-die termination resistors. This field is valid only when POdtOff=0. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>240 ohms +/- 20%</td> </tr> <tr> <td>001b</td> <td>120 ohms +/- 20%</td> </tr> <tr> <td>010b</td> <td>80 ohms +/- 20%</td> </tr> <tr> <td>011b</td> <td>60 ohms +/- 20%</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	240 ohms +/- 20%	001b	120 ohms +/- 20%	010b	80 ohms +/- 20%	011b	60 ohms +/- 20%	111b-100b	Reserved
Bits	Description												
000b	240 ohms +/- 20%												
001b	120 ohms +/- 20%												
010b	80 ohms +/- 20%												
011b	60 ohms +/- 20%												
111b-100b	Reserved												

3	Reserved.												
2:0	DqsDrvStren: DQS drive strength. Read-write. Specifies the drive strength of the DQS pins. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0.75x</td> </tr> <tr> <td>001b</td> <td>1.0x</td> </tr> <tr> <td>010b</td> <td>1.25x</td> </tr> <tr> <td>011b</td> <td>1.5x</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	0.75x	001b	1.0x	010b	1.25x	011b	1.5x	111b-100b	Reserved
Bits	Description												
000b	0.75x												
001b	1.0x												
010b	1.25x												
011b	1.5x												
111b-100b	Reserved												

D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0] Data Byte Transmit PreDriver Calibration 2

Cold reset: xxxx_xxxxh. BIOS: See 2.9.5.4.4.

Table 150: [Index Mapping](#) for D18F2x9C_x0D0F_0[F,7:0]06_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]							
	0706h	0606h	0506h	0406h	0306h	0206h	0106h	0006h
0D0Fh	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 151: [Index Mapping](#) for D18F2x9C_x0D0F_0[F,7:0]0A_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]							
	070Ah	060Ah	050Ah	040Ah	030Ah	020Ah	010Ah	000Ah
0D0Fh	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 152: [Broadcast Mapping](#) for D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]	
	0F0Ah	0F06h
0D0Fh	D18F2x9C_x0D0F_0[7:0]0A	D18F2x9C_x0D0F_0[7:0]06

Bits	Description
31:12	Reserved.
11:6	TxPreP: PMOS predriver calibration code. Read-write. This field specifies the rising edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP] . After updating this value, BIOS must program D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][ValidTxAndPre]=1 for the change to take effect.
5:0	TxPreN: NMOS predriver calibration code. Read-write. This field specifies the falling edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreN] . After updating this value, BIOS must program D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][ValidTxAndPre]=1 for the change to take effect.

D18F2x9C_x0D0F_0[F,7:0]0F_dct[1:0] Data Byte DLL Clock Enable

Cold reset: 0000_0013h.

Table 153: Index addresses for D18F2x9C_x0D0F_0[F,7:0]0F_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]							
	070Fh	060Fh	050Fh	040Fh	030Fh	020Fh	010Fh	000Fh
0D0Fh	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 154: Broadcast write index address for D18F2x9C_x0D0F_0[F,7:0]0F_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F0Fh
0D0Fh	D18F2x9C_x0D0F_0[7:0]0F

Bits	Description								
31:15	Reserved.								
14:12	<p>AlwaysEnDIIclks: always enable DLL clocks. Read-write. BIOS: 2.9.5.4.3.0=DLL clocks are turned off during periods of inactivity. 1=DLL clocks remain on during inactivity. Prior to programming AlwaysEnDIIclks to a value other than 000b, D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][RxMaxDurDIINoLock, TxMaxDurDIINoLock] must both be programmed to 0000b. The bits AlwaysEnDIIclks[2:0] are mapped to DLLs as follows:</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>[2]</td> <td>TxDqs DLL</td> </tr> <tr> <td>[1]</td> <td>TxDq DLL</td> </tr> <tr> <td>[0]</td> <td>RxEn DLL</td> </tr> </tbody> </table>	Bit	Definition	[2]	TxDqs DLL	[1]	TxDq DLL	[0]	RxEn DLL
Bit	Definition								
[2]	TxDqs DLL								
[1]	TxDq DLL								
[0]	RxEn DLL								
11:0	Reserved.								

D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0] Data Byte DLL Power Management

Cold reset: 0000_0000h.

Table 155: Index Addresses for D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_0010h	Byte 0
0D0F_0110h	Byte 1
0D0F_0210h	Byte 2
0D0F_0310h	Byte 3
0D0F_0410h	Byte 4
0D0F_0510h	Byte 5
0D0F_0610h	Byte 6
0D0F_0710h	Byte 7
0D0F_0F10h	Byte [7:0]

Bits	Description
31:13	Reserved.

12	EnRxPadStandby: enable receiver pad standby. Read-write. BIOS: See 2.9.5.7.1=Phy receiver standby mode is enabled to save power when not receiving data. 0=Phy receiver standby mode is disabled.
11:10	Reserved.
9:8	TxDLLWakeupTime: transmit DLL wakeup time. Read-write. BIOS: See 2.9.5.11. Specifies the number of PCLK's that the DLL standby signal must deassert prior to a DLL relock event or before write traffic is sent to transmit DLLs. This field is an alias of D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][TxDLLWakeupTime] .
7	Reserved.
6:4	TxCPUpdPeriod: transmit charge pump DLL wakeup time. Read-write. BIOS: See 2.9.5.11. Specifies the number of DLL relocks required to keep the TxDLLs locked for the period where there is no write traffic. This field is an alias of D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][TxCPUpdPeriod] .
3:0	TxMaxDurDllNoLock: dll no lock. Read-write. BIOS: See 2.9.5.7. This field enables DLL standby and specifies the number of PCLK cycles that occur before the phy DLLs relock. A DLL relock occurs every $\text{MIN}(2^{\text{RxMaxDurDllNoLock}}, 512)$ clock cycles if there are no accesses during the period. 0=Disable the DLL standby power saving feature. This field is an alias of D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][TxMaxDurDllNoLock] .

D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0] Data Byte DLL Power Management

Cold reset: 0000_0000h.

Table 156: Index Addresses for D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_0011h	Byte 0
0D0F_0111h	Byte 1
0D0F_0211h	Byte 2
0D0F_0311h	Byte 3
0D0F_0411h	Byte 4
0D0F_0511h	Byte 5
0D0F_0611h	Byte 6
0D0F_0711h	Byte 7
0D0F_0F11h	Byte [7:0]

Bits	Description
31:10	Reserved.
9:8	RxDLLWakeupTime: transmit DLL wakeup time. Read-write. BIOS: See 2.9.5.11. Specifies the number of PCLK's that the DLL standby signal must deassert prior to a DLL relock event or before read traffic is sent to receive DLLs. This field is an alias of D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][RxDLLWakeupTime] .
7	Reserved.

6:4	RxCPUpdPeriod: transmit charge pump DLL wakeup time. Read-write. BIOS: See 2.9.5.11. Specifies the number of DLL relocks required to keep the RxDLLs locked for the period where there is no read traffic. This field is an alias of D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][RxCPUpdPeriod] .
3:0	RxMaxDurDllNoLock: dll no lock. Read-write. BIOS: See 2.9.5.7. This field enables DLL standby and specifies the number of PCLK cycles that occur before the phy DLLs relock. A DLL relock occurs every $\text{MIN}(2^{\text{RxMaxDurDllNoLock}}, 512)$ clock cycles if there are no accesses during the period. 0=Disable the DLL standby power saving feature. This field is an alias of D18F2x9C_x0000_000D_dct[1:0]_mp[1:0][RxMaxDurDllNoLock] .

D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0] Data Byte DLL Configuration

Table 157: Index Addresses for [D18F2x9C_x0D0F_0\[F,7:0\]13_dct\[1:0\]_mp\[1:0\]](#)

D18F2x98_dct[1:0][31:0]	Function
0D0F_0013h	Byte 0
0D0F_0113h	Byte 1
0D0F_0213h	Byte 2
0D0F_0313h	Byte 3
0D0F_0413h	Byte 4
0D0F_0513h	Byte 5
0D0F_0613h	Byte 6
0D0F_0713h	Byte 7
0D0F_0F13h	Byte [7:0]

Bits	Description
31:15	Reserved.
14	ProcOdtAdv: ProcOdt advance. Read-write. Cold reset: 1. BIOS: $\text{IF}(\text{SODIMM} \ \&\& \ \text{DdrRate} \leq 1333) \ \text{THEN} \ 0 \ \text{ELSE} \ 1 \ \text{ENDIF}$. 0=ProcOdt is asserted 1.5 PCLK before DqsEn . 1=If preceding write, ProcOdt is asserted 2.5 PCLK before DqsEn, else, ProcOdt is asserted 5.0 PCLK before DqsEn.
13:9	Reserved.
8	RxSsbMntClkEn: receive channel maintenance clock enable. Read-write. Cold reset: 0. BIOS: See 2.9.5.11. 1=Enable receive channel maintenance clocks to improve internal timing margin at the cost of some extra power. To enable clock generation, BIOS must first enable maintenance clocks in the DCT (see D18F2x248_dct[1:0]_mp[1:0][RxChMntClkEn]). 0=Disable clocks. The maintenance clock period is fixed at 8 PCLKs when no traffic occurs.
7	RxDqsUDllPowerDown: receive DQS upper DLL power down. Read-write. Cold reset: 0. BIOS: See 2.9.5.11. 1=Power down the upper receiver DQS DLL.
6:2	Reserved.
1	DllDisEarlyU: DLL disable early upper. Read-write. Cold reset: 0. BIOS: See 2.9.5.11. 1=Disable upper receiver DQS DLL early timing for power savings.
0	DllDisEarlyL: DLL disable early lower. Read-write. Cold reset: 0. BIOS: See 2.9.5.11. 1=Disable lower receiver DQS DLL early timing for power savings.

D18F2x9C_x0D0F_0[F,7:0]1C_dct[1:0]_mp[1:0] Data Byte DLL Power Management

Cold reset: 0000_0000h.

Table 158: Index Addresses for D18F2x9C_x0D0F_0[F,7:0]1C_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:8]	Function
0D0F_001Ch	Byte 0
0D0F_011Ch	Byte 1
0D0F_021Ch	Byte 2
0D0F_031Ch	Byte 3
0D0F_041Ch	Byte 4
0D0F_051Ch	Byte 5
0D0F_061Ch	Byte 6
0D0F_071Ch	Byte 7
0D0F_0F1Ch	Byte [7:0]

Bits	Description
15	RxDllStggrEn: Rx DLL stagger enable. Read-write. BIOS: 0. 1=Insert RxDllStggrDly delay prior to waking the receive DLLs from standby. 0=DLLs wake from standby at the default pending-activity reference point.
14	Reserved.
13:8	RxDllStggrDly[5:0]: Rx DLL stagger delay. Read-write. BIOS: 0. Specifies the delay in PCLK cycles from a CAS command or an internal DLL maintenance lock timing event before the receive DLLs exit standby if the data bus is still idle. If data bus traffic is pending and the DLLs are in standby then they wake immediately regardless of the delay value.
7	TxDllStggrEn: Tx DLL stagger enable. Read-write. BIOS: 0. 1=Insert TxDllStggrDly delay prior to waking the transmit DLL from standby. 0=DLL wakes from standby at the default pending-activity reference point.
6	Reserved.
5:0	TxDllStggrDly[5:0]: Tx DLL stagger delay. Read-write. BIOS: 0. Specifies the delay in PCLK cycles from a CAS command or an internal DLL maintenance lock timing event before the transmit DLL exits standby if the data bus is still idle. If data bus traffic is pending and the DLL is in standby then it wakes immediately regardless of the delay value.

D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0] Data Byte Receiver Configuration

Cold reset: 0000_2003h.

Table 159: Index Addresses for D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:8]	Function
0D0F_001Fh	Byte 0
0D0F_011Fh	Byte 1
0D0F_021Fh	Byte 2
0D0F_031Fh	Byte 3
0D0F_041Fh	Byte 4
0D0F_051Fh	Byte 5

Table 159: Index Addresses for D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0]

0D0F_061Fh	Byte 6
0D0F_071Fh	Byte 7
0D0F_0F1Fh	Byte [7:0]

Bits	Description										
31:9	Reserved.										
8	Rx4thStgEn: receiver 4th input stage amplifier enable. Read-write. BIOS: 2.9.5.9.4. 1=Enable 4th stage of amplifiers in the DQ receiver. 0=Disable 4th stage. This bit is ignored if RxBypass3rd4thStg=1.										
7:5	Reserved.										
4:3	RxVioLvl: receiver voltage level. Read-write. BIOS: 2.9.5.4.1. Specifies the VDDIO voltage level. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1.5 V</td> </tr> <tr> <td>01b</td> <td>1.35 V</td> </tr> <tr> <td>10b</td> <td>1.25 V</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	1.5 V	01b	1.35 V	10b	1.25 V	11b	Reserved
Bits	Description										
00b	1.5 V										
01b	1.35 V										
10b	1.25 V										
11b	Reserved										
2	RxBypass3rd4thStg: receiver bypass 3rd and 4th input stage. Read-write. BIOS: 2.9.5.9.4. 1=Bypass the 3rd and 4th stages of amplifiers in the DQ receiver. 0=Disable bypass. This setting adds delay to DQ signals to increase the sampling envelope.										
1:0	Reserved.										

D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0] Data Byte DLL Configuration and Power Down

Cold reset: 0000_0001h.

Table 160: Index Addresses for D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]							
	0730h	0630h	0530h	0430h	0330h	0230h	0130h	0030h
0D0Fh	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 161: Broadcast Write Index Address for D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
0D0Fh	D18F2x9C_x0D0F_0[7:0]30

Bits	Description
31:9	Reserved.
8	BlockRxDqsLock: block rx dqs lock. Read-write. BIOS: 2.9.5.9.3. Specifies how the receive DLLs lock. 1=Lock on PCLK. 0=Lock on both PCLK and the received DQS
7:5	Reserved.
4	PwrDn: power down. Read-write. BIOS: IF (REG==D18F2x9C_x0D0F_0830_dct[1:0] && D18F2x90_dct[1:0][DimmEccEn]==0) THEN 1 ELSE 0 ENDIF. 1=Turn off DLL circuitry.

3	Reserved.
2:0	DllGain: Dll relock gain. Read-write. Specifies the gain of DLL relock circuitry. 000b=Fastest lock time...101b= Slowest lock time. Values of 110b and 111b are reserved.

D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0] Data Byte Fence2 Threshold

BIOS: 2.9.5.4.3.

Table 162: Index addresses for D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]							
	0731h	0631h	0531h	0431h	0331h	0231h	0131h	0031h
0D0Fh	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 163: Broadcast write index address for D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F31h
0D0Fh	D18F2x9C_x0D0F_0[7:0]31

Bits	Description										
31:10	Reserved.										
9	Fence2EnableTxDll: phy fence2 enable transmit DLL. Read-write. Cold reset: 0. 1=Enable the use of Fence2ThresholdTxDll for transmit DLL fence threshold. 0=Use D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxDll] .										
8:5	Fence2ThresholdTxDll: phy fence2 threshold transmit DLL. Read-write. Cold reset: 0. If Fence2EnableTxDll=1, this field specifies the fence delay threshold value used for DQS receiver valid. This field is only used during DQS receiver enable training. See Fence2ThresholdTxPad.										
4	Fence2EnableTxPad: fence2 enable transmit pad. Read-write. Cold reset: 0. 1=Enable the use of Fence2ThresholdTxPad for transmit pad fence threshold. 0=Use D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxPad] .										
3:0	Fence2ThresholdTxPad: phy fence2 threshold transmit pad. Read-write. Cold reset: 0. If Fence2EnableTxPad=1, this field specifies the fence delay threshold value used for write data and write DQS. <table border="0" style="margin-left: 20px;"> <thead> <tr> <th><u>Bits</u></th> <th><u>Definition</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>No delay</td> </tr> <tr> <td>1h</td> <td>1/64 MEMCLK delay</td> </tr> <tr> <td>Eh-2h</td> <td><Fence2ThresholdTxPad> MEMCLK delay</td> </tr> <tr> <td>Fh</td> <td>15/64 MEMCLK delay</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Definition</u>	0h	No delay	1h	1/64 MEMCLK delay	Eh-2h	<Fence2ThresholdTxPad> MEMCLK delay	Fh	15/64 MEMCLK delay
<u>Bits</u>	<u>Definition</u>										
0h	No delay										
1h	1/64 MEMCLK delay										
Eh-2h	<Fence2ThresholdTxPad> MEMCLK delay										
Fh	15/64 MEMCLK delay										

D18F2x9C_x0D0F_2[F,2:0]00_dct[1:0]_mp[1:0] Clock Pad Configuration

Cold reset: xxxx_xxxxh. BIOS: See [D18F2x9C_x0000_0000_dct\[1:0\]_mp\[1:0\]](#). See 2.9.2.1 for clock chip to pin mapping.

Table 164: Index Addresses for D18F2x9C_x0D0F_2[F,2:0]00_dct[1:0]_mp[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_2000h	Clk chip 0
0D0F_2100h	Clk chip 1
0D0F_2200h	Reserved
0D0F_2F00h	All clk chips

Bits	Description												
31:3	Reserved.												
2:0	ClkDrvStren: MEMCLK drive strength. Read-write. This field specifies the drive strength of the MEMCLK pins. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>1.0x</td> </tr> <tr> <td>001b</td> <td>1.25x</td> </tr> <tr> <td>010b</td> <td>1.5x</td> </tr> <tr> <td>011b</td> <td>2.0x</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	1.0x	001b	1.25x	010b	1.5x	011b	2.0x	111b-100b	Reserved
Bits	Description												
000b	1.0x												
001b	1.25x												
010b	1.5x												
011b	2.0x												
111b-100b	Reserved												

D18F2x9C_x0D0F_2[2:0]02_dct[1:0] Clock Transmit PreDriver Calibration

Cold reset: xxxx_xxxxh. BIOS: See 2.9.5.4.4.

Table 165: Index Mapping for D18F2x9C_x0D0F_2[2:0]02_dct[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_2002h	Clock 0 Pad Group 0
0D0F_2102h	Clock 1 Pad Group 0
0D0F_2202h	Clock 2 Pad Group 0

Bits	Description
31:16	Reserved.
15	ValidTxAndPre: predriver calibration code valid. Read-write; cleared-by-hardware. 1=Predriver calibration codes are copied from this register into the associated transmit pad. Hardware clears this field after the copy is complete.
14:12	Reserved.
11:6	TxPreP: PMOS predriver calibration code. Read-write. This field specifies the rising edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP]. After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.
5:0	TxPreN: NMOS predriver calibration code. Read-write. This field specifies the falling edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreN]. After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.

D18F2x9C_x0D0F_[C,8,2][2:0]1F Receiver Configuration

Cold reset: 0000_2000h.

Table 166: Index address mapping for D18F2x9C_x0D0F_[C,8,2][2:0]1F

D18F2x98_dct[1:0][31:0]	Function
0D0F_201Fh	Clock 0
0D0F_211Fh	Clock 1
0D0F_221Fh	Clock 2
0D0F_801Fh	Cmd/Addr 0
0D0F_811Fh	Cmd/Addr 1
0D0F_821Fh	Reserved
0D0F_C01Fh	Address
0D0F_C11Fh	Reserved
0D0F_C21Fh	Reserved

Bits	Description										
31:5	Reserved.										
4:3	RxVioLvl: receiver voltage level. Read-write. BIOS: See 2.9.5.4.1. Specifies the VDDIO voltage level. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1.5 V</td> </tr> <tr> <td>01b</td> <td>1.35 V</td> </tr> <tr> <td>10b</td> <td>1.25 V</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	1.5 V	01b	1.35 V	10b	1.25 V	11b	Reserved
Bits	Description										
00b	1.5 V										
01b	1.35 V										
10b	1.25 V										
11b	Reserved										
2:0	Reserved.										

D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0] Clock DLL Configuration and Power Down

Cold reset: 0000_0001h.

Table 167: Index Addresses for D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_2030h	Clk chip 0
0D0F_2130h	Clk chip 1
0D0F_2230h	Reserved
0D0F_2F30h	All clk chips

Bits	Description
31:16	Reserved.
15:5	Reserved.
4	PwrDn: power down. Read-write. 1=Turn off DLL circuitry.
3	Reserved.
2:0	DllGain: Dll relock gain. Read-write. See D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0][DllGain] for a description.

D18F2x9C_x0D0F_4003_dct[1] DCT0 M1 Fence Value

Cold reset: 0000_3F3Fh. This register is used by BIOS to temporarily store the trained fence values for DCT0 for the M1 memory P-state, and has no effect on the operation of the hardware. See 2.9.5.4.3. This register is accessed through DCT1 only, and is available regardless of D18F2x94_dct[1:0][DisDramInterface].

Bits	Description
31:15	Reserved.
14:10	TxDll: Transmit dll fence value. Read-write. BIOS saves the trained TxDll to this field until such time as it is able to write the value to NVRAM.
9:5	RxDll: Receiver dll fence value. Read-write. BIOS saves the trained RxDll to this field until such time as it is able to write the value to NVRAM.
4:0	TxPad: Transmit pad fence value. Read-write. BIOS saves the trained TxPad to this field until such time as it is able to write the value to NVRAM.

D18F2x9C_x0D0F_4004_dct[1] DCT1 M1 Fence Value

Cold reset: 0000_0000h. This register is used by BIOS to temporarily store the trained fence values for DCT1 for the M1 memory P-state, and has no effect on the operation of the hardware. See 2.9.5.4.3. This register is accessed through DCT1 only, and is available regardless of D18F2x94_dct[1:0][DisDramInterface].

Bits	Description
31:15	Reserved.
14:10	TxDll: Transmit dll fence value. Read-write. BIOS saves the trained TxDll to this field until such time as it is able to write the value to NVRAM.
9:5	RxDll: Receiver dll fence value. Read-write. BIOS saves the trained RxDll to this field until such time as it is able to write the value to NVRAM.
4:0	TxPad: Transmit pad fence value. Read-write. BIOS saves the trained TxPad to this field until such time as it is able to write the value to NVRAM.

D18F2x9C_x0D0F_4009_dct[1:0] Phy Cmp Configuration

Cold reset: 0000_2000h.

Bits	Description										
31:16	Reserved.										
15:14	CmpVioLvl: cmp voltage level. Read-write. BIOS: See 2.9.5.4.1. This field specifies the VDDIO voltage level. Setting this field in DCT0 adjusts the VDDIO voltage level for DCT0 and DCT1. Setting this field in DCT1 has no effect. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1.5 V</td> </tr> <tr> <td>01b</td> <td>1.35 V</td> </tr> <tr> <td>10b</td> <td>1.25 V</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	1.5 V	01b	1.35 V	10b	1.25 V	11b	Reserved
Bits	Description										
00b	1.5 V										
01b	1.35 V										
10b	1.25 V										
11b	Reserved										
13:4	Reserved.										

3:2	ComparatorAdjust: comparator adjust. Read-write. BIOS: D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0][RxVioLvl] . This field specifies the adjustment signals for the comparator differential amplifier. Setting this field in DCT0 adjusts the comparator for DCT0 and DCT1. Setting this field in DCT1 has no effect.
1:0	Reserved.

D18F2x9C_x0D0F_8[1:0]0[8,4,0]_dct[1:0]_mp[1:0] CAD1 Pad Configuration

Cold reset: xxxx_xxxxh. BIOS: See [D18F2x9C_x0000_0000_dct\[1:0\]_mp\[1:0\]](#). See 2.9.2.1 for chip and group to pin mapping.

Table 168: Index Mapping for [D18F2x9C_x0D0F_8\[1:0\]0\[8,4,0\]_dct\[1:0\]_mp\[1:0\]](#)

Address Bits	Valid Values	Name
D18F2x98_dct[1:0][3:2]	[2:0]	Pad Group
D18F2x98_dct[1:0][6]	[1:0]	Memory Pstate
D18F2x98_dct[1:0][11:8]	[1:0]	Cmd/Addr chip
D18F2x98_dct[1:0][11:8]	Fh	Cmd/Addr chip[1:0]

Bits	Description												
31:3	Reserved.												
2:0	DrvStrength: address/command/csodt drive strength. Read-write. This field specifies the drive strength of the address, CS, ODT, RAS, CAS, WE, bank and parity pins. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>1.0x</td> </tr> <tr> <td>001b</td> <td>1.25x</td> </tr> <tr> <td>010b</td> <td>1.5x</td> </tr> <tr> <td>011b</td> <td>2.0x</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	1.0x	001b	1.25x	010b	1.5x	011b	2.0x	111b-100b	Reserved
Bits	Description												
000b	1.0x												
001b	1.25x												
010b	1.5x												
011b	2.0x												
111b-100b	Reserved												

D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0] Transmit PreDriver Calibration

Cold reset: xxxx_xxxxh. BIOS: See 2.9.5.4.4.

Table 169: Index Mapping for [D18F2x9C_x0D0F_\[C,8\]\[1:0\]02_dct\[1:0\]](#)

D18F2x98_dct[1:0][31:0]	Function
0D0F_8002h	Cmd/Addr 0 Pad Group 0
0D0F_8102h	Cmd/Addr 1 Pad Group 0
0D0F_C002h	Address Pad Group 0

Bits	Description
31:16	Reserved.
15	ValidTxAndPre: predriver calibration code valid. Read-write; cleared-by-hardware. 1=Predriver calibration codes are copied from this register and D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0] into the associated transmit pad. Hardware clears this field after the copy is complete.
14:12	Reserved.

11:6	TxPreP: PMOS predriver calibration code. Read-write. This field specifies the rising edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP] . After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.
5:0	TxPreN: NMOS predriver calibration code. Read-write. This field specifies the falling edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreN] . After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.

D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0] Transmit PreDriver Calibration 2

Cold reset: xxxx_xxxxh. BIOS: See 2.9.5.4.4.

Table 170: Index Mapping for [D18F2x9C_x0D0F_\[C,8\]\[1:0\]\[12,0E,0A,06\]_dct\[1:0\]](#)

D18F2x98_dct[1:0][31:0]	Function
0D0F_8006h	Cmd/Addr 0 Pad Group 1
0D0F_800Ah	Cmd/Addr 0 Pad Group 2
0D0F_8106h	Cmd/Addr 1 Pad Group 1
0D0F_810Ah	Cmd/Addr 1 Pad Group 2
0D0F_C006h	Address Pad Group 1
0D0F_C00Ah	Address Pad Group 2
0D0F_C00Eh	Address Pad Group 3
0D0F_C012h	Address Pad Group 4

Bits	Description
31:12	Reserved.
11:6	TxPreP: PMOS predriver calibration code. Read-write. This field specifies the rising edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreP] . After updating this value, BIOS must program D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0][ValidTxAndPre]=1 for the change to take effect.
5:0	TxPreN: NMOS predriver calibration code. Read-write. This field specifies the falling edge slew rate of the transmit pad. See: D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0][TxPreN] . After updating this value, BIOS must program D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0][ValidTxAndPre]=1 for the change to take effect.

D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0] DLL Delay and Configuration

Cold reset: 0000_8000h. BIOS: [D18F2x9C_x0000_0004_dct\[1:0\]_mp\[1:0\]](#). See 2.9.2.1 for chip to pin mapping.

Table 171: Index Mapping for [D18F2x9C_x0D0F_8\[1:0\]20_dct\[1:0\]_mp\[1:0\]](#)

Address Bits	Valid Values	Name
D18F2x98_dct[1:0][6]	[1:0]	Memory Pstate
D18F2x98_dct[1:0][11:8]	[1:0]	Cmd/Addr 1 chip
D18F2x98_dct[1:0][11:8]	Fh	Cmd/Addr 1 chip[1:0]
Note: Control of certain pins in Cmd/Addr 0 chip is overridden with D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][DiffTimingEn]		

Bits	Description								
31:8	Reserved.								
7	FenceBit: fence bit. Read-write; updated-by-hardware. 1= A 0.5-PCLK-delayed version of the signal from the TxFIFO is used for the input of the output pad. 0= An undelayed version is used for the input of the output pad.								
6	Fence2Bit: second fence bit. Read-write; updated-by-hardware. 1= A 1-PCLK-delayed version of the signal from the TxFIFO is used for the input of the output pad. FenceBit should be programmed 1b. 0= FenceBit determines the input to the output pad.								
5	Reserved.								
4:0	Delay: delay. Read-write. Specifies the time that the signals are delayed from the default setup time. <table border="1" data-bbox="267 646 868 787"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0/64 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><Delay>/64 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	00h	0/64 MEMCLK delay	1Eh-01h	<Delay>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay
Bits	Description								
00h	0/64 MEMCLK delay								
1Eh-01h	<Delay>/64 MEMCLK delay								
1Fh	31/64 MEMCLK delay								

D18F2x9C_x0D0F_8021_dct[1:0]_mp[1:0] DLL CS 6 & 7 Timing Control

Cold reset: 0000_0000h. BIOS: [2.9.5.4.3.D18F2x9C_x0000_0004_dct\[1:0\]_mp\[1:0\]](#) must be reprogrammed any time new values are written to this register.

Bits	Description								
31:16	Reserved.								
15	DiffTimingEn: differential timing enable. Read-write. Enables independent timing controls for CS[7:6]. 1=Timing is specified by Fence and Delay. 0=Timing is specified by D18F2x9C_x0000_0004_dct[1:0]_mp[1:0][CsOdtFineDelay] .								
14:8	Reserved.								
7	Fence: fence. Read-write. Adjusts the phase relationship between the fifo and the pad when DiffTimingEn=1.								
6:5	Reserved.								
4:0	Delay: delay. Read-write. Specifies the time that the CS6 & 7 pins are delayed from the default setup time when DiffTimingEn=1. <table border="1" data-bbox="267 1434 868 1570"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0/64 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><Delay>/64 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	00h	0/64 MEMCLK delay	1Eh-01h	<Delay>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay
Bits	Description								
00h	0/64 MEMCLK delay								
1Eh-01h	<Delay>/64 MEMCLK delay								
1Fh	31/64 MEMCLK delay								

D18F2x9C_x0D0F_812F_dct[1:0] Tristate Configuration

Reset: 0000_00A0h. BIOS: See [2.9.5.11](#).

Bits	Description
31:1	Reserved.
0	PARTri: MEMPAR tri-state. Read-write. Specifies tri-state control for the memory parity signal. 1=Signal is tri-stated. 0=Signal is not tri-stated.

D18F2x9C_x0D0F_[C,8][F:0]30_dct[1:0] Cmd/Addr DLL Configuration and Power Down

Cold reset: 0000_0001h. BIOS: See 2.9.5.4.4. See Table 19 and Table 20 for chip to pin mapping.

Table 172: Index Mapping for D18F2x9C_x0D0F_[C,8][F:0]30_dct[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_8030h	Cmd/Addr 0
0D0F_8130h	Cmd/Addr 1
0D0F_8F30h	Cmd/Addr 0 and Cmd/Addr 1
0D0F_C030h	Address 2

Bits	Description
15:5	Reserved.
4	PwrDn: power down. Read-write. 1=Turn off DLL circuitry.
3	Reserved.
2:0	DllGain: Dll relock gain. Read-write. See D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0][DllGain] for a description.

D18F2x9C_x0D0F_C000_dct[1:0]_mp[1:0] CKE 2.0X Pad Configuration

Reset: 0000_0003h. BIOS: See 2.9.5.11.

Bits	Description												
31:9	Reserved.												
8	LowPowerDrvStrengthEn: low power drive strength enable. Read-write. 1=CKE driver is forced to low drive strength when phy is in self-refresh mode. 0=CKE is driven with strength specified by D18F2x9C_x0000_0000_dct[1:0]_mp[1:0][CkeDrvStren] when phy is in self-refresh mode. This bit may only be set if CkeDrvStren==010b or 011b.												
7:3	Reserved.												
2:0	CkeDrvStren: CKE drive strength. Read-write. This field specifies the drive strength of the CKE pins. <table border="1" data-bbox="267 1375 617 1585"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>1.0x</td> </tr> <tr> <td>001b</td> <td>1.25x</td> </tr> <tr> <td>010b</td> <td>1.5x</td> </tr> <tr> <td>011b</td> <td>2.0x</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	1.0x	001b	1.25x	010b	1.5x	011b	2.0x	111b-100b	Reserved
Bits	Description												
000b	1.0x												
001b	1.25x												
010b	1.5x												
011b	2.0x												
111b-100b	Reserved												

D18F2x9C_x0D0F_C0[10,0C,08,04]_dct[1:0]_mp[1:0] CAD2 Pad Configuration

Cold reset: xxxx_xxxxh. BIOS: See [D18F2x9C_x0000_0000_dct\[1:0\]_mp\[1:0\]](#). See 2.9.2.1 for Address2 chip group to pin mapping.

Table 173: Index Mapping for [D18F2x9C_x0D0F_C0\[10,0C,08,04\]_dct\[1:0\]_mp\[1:0\]](#)

Address Bits	Valid Values	Name
D18F2x98_dct[1:0][4:2]	[4:1]	Pad Group
D18F2x98_dct[1:0][6]	[1:0]	Memory Pstate
D18F2x98_dct[1:0][15:12]	Ch	Address2 chip

Bits	Description												
31:3	Reserved.												
2:0	DrvStrength: address drive strength. Read-write. Cold reset: 011b. This field specifies the drive strength of the address, CS, ODT, RAS, CAS, WE, bank and parity pins. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>1.0x</td> </tr> <tr> <td>001b</td> <td>1.25x</td> </tr> <tr> <td>010b</td> <td>1.5x</td> </tr> <tr> <td>011b</td> <td>2.0x</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	1.0x	001b	1.25x	010b	1.5x	011b	2.0x	111b-100b	Reserved
Bits	Description												
000b	1.0x												
001b	1.25x												
010b	1.5x												
011b	2.0x												
111b-100b	Reserved												

D18F2x9C_x0D0F_C020_dct[1:0]_mp[1:0] CAD2 Address DLL Delay and Configuration

Cold reset: 0000_8000h. BIOS: [D18F2x9C_x0000_0004_dct\[1:0\]_mp\[1:0\]](#). This register controls timing for signals on Address[15:0] pins.

Table 174: Index Mapping for [D18F2x9C_x0D0F_C020_dct\[1:0\]_mp\[1:0\]](#)

Address Bits	Valid Values	Name
D18F2x98_dct[1:0][6]	[1:0]	Memory Pstate

Bits	Description
31:8	Reserved.
7	FenceBit: fence bit. See: D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][FenceBit] .
6	Fence2Bit: second fence bit. See: D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Fence2Bit] .
5	Reserved.
4:0	Delay: delay. See: D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Delay] .

D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0] CAD2 CKE DLL Delay and Configuration

Cold reset: 0000_0000h. BIOS: [D18F2x9C_x0000_0004_dct\[1:0\]_mp\[1:0\]](#). This register controls timing for signals on CKE[3:0] pads. See 2.9.2.1 for pad to pin mapping.

Table 175: Index Mapping for D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0]

Address Bits	Valid Values	Name
D18F2x98_dct[1:0][6]	[1:0]	Memory Pstate

Bits	Description
31:8	Reserved.
7	FenceBit: fence bit. See: D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][FenceBit].
6	Fence2Bit: second fence bit. See: D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Fence2Bit].
5	Reserved.
4:0	Delay: delay. See: D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0][Delay].

D18F2x9C_x0D0F_E000_dct[1:0]_mp[1:0] Phy Master Configuration

Cold reset: 0000_0002h.

Bits	Description
31:16	Reserved.
15:14	Rate[4:3]. See: Rate[2:0].
13:4	Reserved.
3	FreqValid: Frequency valid. Read-write. 1=The memory frequency specified with the Rate field is valid. software must set this bit when programming Rate for the setting to take effect. Regardless of this bit, a new rate does not take effect unless the target register context is of the current memory P-state, or until the memory P-state is changed to the target Rate context.
2:0	Rate[2:0]. Read-write. Rate[4:0] = {Rate[4:3], Rate[2:0]}. Rate[4:0] is the 5-bit DDR data rate that the phy uses to configure the PLL. See Table 132 [Memory Clock Frequency Value Definition] . Software must set FreqValid=1 when programming this field. If memory P-states are enabled then software must program the M1 context. It is not necessary for software to program the M0 context when using D18F2x94_dct[1:0][MemClkFreqVal] .

D18F2x9C_x0D0F_E006_dct[1:0] Phy PLL Lock Time

Cold reset: 0000_0190h.

Bits	Description
31:16	Reserved.
15:0	PlILockTime: pll lock time. Read-write. Specifies the number of 5ns periods the phy waits for PLLs to lock during a frequency change. See 2.9.5.4.2 [DRAM Channel Frequency Change] .

D18F2x9C_x0D00_E008_dct[1:0] Phy Master Configuration

Cold reset: 0000_0013h.

Bits	Description
31:13	Reserved.

12	PhyPS: Phy memory P-state. Read-only; updated-by-hardware. Indicates the current memory P-state for the phy. 0=M0, 1=M1.
11:5	Reserved.
4:0	FenceValue: fence value. Read-only; updated-by-hardware. Indicates the fence value used to create the fence bit in the DLL delay registers. See 2.9.5.4.3 [Phy Fence Programming] .

D18F2x9C_x0D04_E008_dct[1:0] Phy Master Configuration

Cold reset: 0000_0013h.

Bits	Description
31:9	Reserved.
8	PStateToAccess: P-state to access. Read-write. Specifies the memory P-state context for context sensitive DDR phy CSR accesses using D18F2x98_dct[1:0] when <code>DctOffset[29:20] != 0D0h</code> . 0=M0, 1=M1. The processor updates this field during a memory P-state change with the target state. This field is only valid in the master channel. See D18F2x9C_x0D0F_E018_dct[1:0] [PhyPSMasterChannel]. See D18F2x94_dct[1:0] [DphyMemPsSelEn].
7:0	Reserved.

D18F2x9C_x0D0F_E00A_dct[1:0] Phy Dynamic Power Mode

Cold reset: 0000_0000h.

Bits	Description
31:5	Reserved.
4	SkewMemClk: skew MEMCLK. Read-write. BIOS: IF ((REG==D18F2x9C_x0D0F_E00A_dct[0]) & ~D18F2x94_dct[0][DisDramInterface] & ~D18F2x94_dct[1][DisDramInterface]) THEN 1 ELSE 0 ENDIF. 1=Skew MEMCLK signals with respect to other channel. SkewMemClk must be 0 if D18F2x94_dct[1:0][DisDramInterface]=1 for either channel. This bit must be set prior to setting D18F2x94_dct[1:0][MemClkFreqVal] during DRAM initialization. See 2.9.5.4.2 [DRAM Channel Frequency Change] .
3:0	Reserved.

D18F2x9C_x0D0F_E010_dct[1:0] DLL Gain Sequencer Timer 0

Cold reset: 0000_00C0h.

The gain sequencer allows multiple gain settings during reset/relock of DLLs. Nominally, the DLLs start with the highest gain and remains in that mode for `EnDllGain0` amount of time. The sequencer then selects the next highest gain and so on until all gain settings are exhausted, remaining in a low gain tracking mode. A thresholding register allows further tailoring for frequency-independent setup of the sequencer register. .

Bits	Description
31:15	Reserved.
14:10	EnDllGain0[4:0]: Enable Dll gain G=0. Read-write. BIOS: 2.9.5.4. Specifies the number of PCLK cycles times 16 that the Dll gain is configured to a value of “G” during a relock event. See D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0][DllGain] for a definition of gain values.
9:5	EnDllGain1[4:0]: Enable Dll gain G=1. Read-write. BIOS: 2.9.5.4. See EnDllGain0.
4:0	EnDllGain2[4:0]: Enable Dll gain G=2. Read-write. BIOS: 2.9.5.4. See EnDllGain0.

D18F2x9C_x0D0F_E011_dct[1:0] DLL Gain Sequencer Timer 1

Cold reset: 0000_0020h.

Bits	Description
31:15	Reserved.
14:10	EnDllGain3[4:0]: Enable Dll gain G=3. Read-write. BIOS: 2.9.5.4. See D18F2x9C_x0D0F_E010_dct[1:0][EnDllGain0] .
9:5	EnDllGain4[4:0]: Enable Dll gain G=4. Read-write. BIOS: 2.9.5.4. See D18F2x9C_x0D0F_E010_dct[1:0][EnDllGain0] .
4:0	EnDllGain5[4:0]: Enable Dll gain G=5. Read-write. BIOS: 2.9.5.4. See D18F2x9C_x0D0F_E010_dct[1:0][EnDllGain0] .

D18F2x9C_x0D0F_E012_dct[1:0] DLL Gain Frequency Threshold

Cold reset: 0000_33FFh.

Bits	Description
31:15	Reserved.
14:10	DllGainFreqThresholdG0[4:0]: Dll gain G=0 frequency threshold. Read-write. BIOS: 2.9.5.4. Specifies a minimum FreqCode[4:0] threshold to skip gain = G in the relock sequence.
9:5	DllGainFreqThresholdG1[4:0]: Dll gain G=1 frequency threshold. Read-write. BIOS: 2.9.5.4. See DllGainFreqThresholdG0.
4:0	DllGainFreqThresholdG2[4:0]: Dll gain G=2 frequency threshold. Read-write. BIOS: 2.9.5.4. See DllGainFreqThresholdG0.

D18F2x9C_x0D0F_E013_dct[1:0] Phy PLL Regulator Wait Time

Cold reset: 0000_00D8h.

Bits	Description
31:16	Reserved.
15:0	PllRegWaitTime: PLL regulator wait time. Read-write. BIOS: 2.9.5.4 . 1=Specifies the number of 5 ns periods the phy waits for the PLL to become stable when coming out of PLL regulator off power down mode.

D18F2x9C_x0D0F_E018_dct[1:0] Fence

Cold reset: 0000_0013h. BIOS: [2.9.5.4.3](#).

Bits	Description
31:9	Reserved.
8	PhyPSMasterChannel: phy memory P-state master channel. Read-write. 0=Channel 0 updates the PLL during a memory P-state change. 1=Channel 1 updates the PLL during a memory P-state change. BIOS should set this bit only if channel 0 has no memory and is disabled. This field is valid for D18F2x9C_x0D0F_E018_dct[0] only.
7:5	Reserved.
4:0	Fence: fence value. Read-write; updated-by-hardware. This field specifies the fence delay threshold value used for command, address, write data and write DQS.

D18F2x9C_x0D0F_E019_dct[1:0] Fence2

Cold reset: 0000_0000h. BIOS: [2.9.5.4.3](#).

Bits	Description										
31:15	Reserved.										
14	Fence2EnableRxDll: phy fence2 enable receive DLL. Read-write. 1=Enable the use of Fence2ThresholdRxDll for DQS receiver enable second fence threshold. 0=Disable second fence. See D18F2x9C_x0000_000C_dct[1:0][FenceThresholdRxDll] for first fence.										
13:10	Fence2ThresholdRxDll: phy fence2 threshold DQS receiver enable. Read-write. If Fence2EnableRxDll=1, this field specifies the fence delay threshold value used for DQS receiver enable. See Fence2ThresholdTxPad.										
9:5	Reserved.										
4	Fence2EnableTxPad: fence2 enable transmit pad. Read-write. 1=Enable the use of Fence2ThresholdTxPad for transmit pad second fence threshold. 0=Disable second fence. See D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxPad] for first fence.										
3:0	Fence2ThresholdTxPad: phy fence2 threshold transmit pad. Read-write. If Fence2EnableTxPad=1, this field specifies the second fence delay threshold value used for command, address, write data and write DQS. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>16/64 MEMCLK delay</td> </tr> <tr> <td>1h</td> <td>17/64 MEMCLK delay</td> </tr> <tr> <td>Eh-2h</td> <td>{1,<Fence2ThresholdTxPad>}/64 MEMCLK delay</td> </tr> <tr> <td>Fh</td> <td>31/64 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Definition	0h	16/64 MEMCLK delay	1h	17/64 MEMCLK delay	Eh-2h	{1,<Fence2ThresholdTxPad>}/64 MEMCLK delay	Fh	31/64 MEMCLK delay
Bits	Definition										
0h	16/64 MEMCLK delay										
1h	17/64 MEMCLK delay										
Eh-2h	{1,<Fence2ThresholdTxPad>}/64 MEMCLK delay										
Fh	31/64 MEMCLK delay										

D18F2xA4 DRAM Controller Temperature Throttle

See 2.9.1 [DCT Configuration Registers] and 2.9.9 [DRAM On DIMM Thermal Management and Power Capping].

Bits	Description																																		
31:24	Reserved.																																		
23:20	<p>BwCapCmdThrottleMode: bandwidth capping command throttle mode. Read-write. Reset: 0. Specifies the command throttle mode when BwCapEn=1. The DCT throttles commands over a rolling window of 100 clock cycles, maintaining the average throttling as specified by this field.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>Command throttling is disabled</td> </tr> <tr> <td>0001b</td> <td>Throttle commands by 30%</td> </tr> <tr> <td>0010b</td> <td>Throttle commands by 40%</td> </tr> <tr> <td>0011b</td> <td>Throttle commands by 50%</td> </tr> <tr> <td>0100b</td> <td>Throttle commands by 55%</td> </tr> <tr> <td>0101b</td> <td>Throttle commands by 60%</td> </tr> <tr> <td>0110b</td> <td>Throttle commands by 65%</td> </tr> <tr> <td>0111b</td> <td>Throttle commands by 70%</td> </tr> <tr> <td>1000b</td> <td>Throttle commands by 75%</td> </tr> <tr> <td>1001b</td> <td>Throttle commands by 80%</td> </tr> <tr> <td>1010b</td> <td>Throttle commands by 85%</td> </tr> <tr> <td>1011b</td> <td>Throttle commands by 90%</td> </tr> <tr> <td>1100b</td> <td>Throttle commands by 95%</td> </tr> <tr> <td>1101b</td> <td>Reserved</td> </tr> <tr> <td>1110b</td> <td>Throttle commands as specified by CmdThrottleMode</td> </tr> <tr> <td>1111b</td> <td>Reserved</td> </tr> </tbody> </table> <p>Throttling should not be enabled until after DRAM initialization (D18F2x110[DramEnable]=1) and training (see 2.9.5.9 [DRAM Training]) are complete.</p>	Bits	Description	0000b	Command throttling is disabled	0001b	Throttle commands by 30%	0010b	Throttle commands by 40%	0011b	Throttle commands by 50%	0100b	Throttle commands by 55%	0101b	Throttle commands by 60%	0110b	Throttle commands by 65%	0111b	Throttle commands by 70%	1000b	Throttle commands by 75%	1001b	Throttle commands by 80%	1010b	Throttle commands by 85%	1011b	Throttle commands by 90%	1100b	Throttle commands by 95%	1101b	Reserved	1110b	Throttle commands as specified by CmdThrottleMode	1111b	Reserved
Bits	Description																																		
0000b	Command throttling is disabled																																		
0001b	Throttle commands by 30%																																		
0010b	Throttle commands by 40%																																		
0011b	Throttle commands by 50%																																		
0100b	Throttle commands by 55%																																		
0101b	Throttle commands by 60%																																		
0110b	Throttle commands by 65%																																		
0111b	Throttle commands by 70%																																		
1000b	Throttle commands by 75%																																		
1001b	Throttle commands by 80%																																		
1010b	Throttle commands by 85%																																		
1011b	Throttle commands by 90%																																		
1100b	Throttle commands by 95%																																		
1101b	Reserved																																		
1110b	Throttle commands as specified by CmdThrottleMode																																		
1111b	Reserved																																		
19:15	Reserved.																																		
14:12	<p>CmdThrottleMode: command throttle mode. Read-write. Reset: 0. BIOS: See 2.9.5.7. Specifies the command throttle mode when ODTSEn=1 and the EVENT_L pin is asserted. The DCT throttles commands over a rolling window of 100 clock cycles, maintaining the average throttling as specified by this field.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Command throttling is disabled.</td> </tr> <tr> <td>001b</td> <td>Throttle commands by 30%.</td> </tr> <tr> <td>010b</td> <td>Throttle commands by 50%.</td> </tr> <tr> <td>011b</td> <td>Throttle commands by 60%.</td> </tr> <tr> <td>100b</td> <td>Throttle commands by 70%.</td> </tr> <tr> <td>101b</td> <td>Throttle commands by 80%.</td> </tr> <tr> <td>110b</td> <td>Throttle commands by 90%.</td> </tr> <tr> <td>111b</td> <td>Place the DRAM devices in powerdown mode (see D18F2x94_dct[1:0][PowerDownMode]) when EVENT_L is asserted. This mode is not valid if D18F2x94_dct[1:0][PowerDownEn]=0.</td> </tr> </tbody> </table> <p>Throttling should not be enabled until after DRAM initialization (D18F2x110[DramEnable]=1) and training (see 2.9.5.9 [DRAM Training]) are complete. See also BwCapEn.</p>	Bits	Description	000b	Command throttling is disabled.	001b	Throttle commands by 30%.	010b	Throttle commands by 50%.	011b	Throttle commands by 60%.	100b	Throttle commands by 70%.	101b	Throttle commands by 80%.	110b	Throttle commands by 90%.	111b	Place the DRAM devices in powerdown mode (see D18F2x94_dct[1:0][PowerDownMode]) when EVENT_L is asserted. This mode is not valid if D18F2x94_dct[1:0][PowerDownEn]=0 .																
Bits	Description																																		
000b	Command throttling is disabled.																																		
001b	Throttle commands by 30%.																																		
010b	Throttle commands by 50%.																																		
011b	Throttle commands by 60%.																																		
100b	Throttle commands by 70%.																																		
101b	Throttle commands by 80%.																																		
110b	Throttle commands by 90%.																																		
111b	Place the DRAM devices in powerdown mode (see D18F2x94_dct[1:0][PowerDownMode]) when EVENT_L is asserted. This mode is not valid if D18F2x94_dct[1:0][PowerDownEn]=0 .																																		

11	BwCapEn: bandwidth capping enable. Read-write. Reset: 0. 1=The memory command throttle mode specified by BwCapCmdThrottleMode is applied. This bit is used by software to enable command throttling independent of the state of the EVENT_L pin. If this bit is set, ODTSEn=1, and the EVENT_L pin is asserted, the larger of the two throttle percentages specified by CmdThrottleMode and BwCapCmdThrottleMode is used.
10:9	Reserved.
8	ODTSEn: on DIMM temperature sensor enable. Read-write. Reset: 0. BIOS: See 2.9.5.7. Enables the monitoring of the EVENT_L pin and indicates whether the on DIMM temperature sensors of the DIMMs on a channel are enabled. 0=Disabled. 1=Enabled. While the EVENT_L pin is asserted, the controller (a) doubles the refresh rate (if Tref=7.8 us), and (b) throttles the address bus utilization as specified by CmdThrottleMode[2:0].
7:0	Reserved.

D18F2xA8_dct[1:0] DRAM Controller Miscellaneous 2

See 2.9.1 [DCT Configuration Registers].

Bits	Description
31	PerRankTimingEn: per rank timing enable. Read-write. Reset: 0. BIOS: 1. Specifies the mapping between chip selects and a set of programmable timing delays. 1=Each chip select is controlled by a set of timing delays. 0=Each chip select pair is controlled by a set timing delays. See D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0] , D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0] , D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0] , and D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0] .
30	Reserved.
29	RefChCmdMgtDis: refresh channel command management disable. Read-write. Reset: 0. 1=DCTs issue refresh commands independently. 0=DCTs stagger the issue of refresh commands. This bit must be set the same in all DCTs.
28	FastSelfRefEntryDis: fast self refresh entry disable. Read-write. Reset: 1. BIOS: ~ D18F2x1B4[FlushWrOnS3StpGnt] . 1=DCT pushes outstanding transactions to memory prior to entering self refresh. 0=DCT enters self refresh immediately unless instructed to push outstanding transactions to memory by D18F2x11C[FlushWrOnStpGnt] or D18F2x1B4[FlushWrOnS3StpGnt] . This bit must be programmed to the same value on all enabled DCTs.
27	CsMux67: chip select mux 6 and 7. See: CsMux45. See also CSTimingMux67.
26	CsMux45: chip select mux 4 and 5. Read-Write. Reset: 0. Specifies the DCT mode used to drive the associated chip selects. 0=Chip select mode. 1=Extended address mode; A[17:16] are driven on the DRAM CS pins.
25:24	Reserved.
23	Reserved.
22	PrtlChPDEnhEn: partial channel power down enh enable. Read-write. Reset: 0. BIOS: 0. Selects the channel idle hysteresis for fast exit/slow exit mode changes when (D18F2x94_dct[1:0][Power-DownMode] & D18F2x84_dct[1:0][PchgPDMoSel]). 1=Hysteresis specified by D18F2x244_dct[1:0][PrtlChPDDynDly] . 0=256 clock hysteresis.
21	AggrPDEn: aggressive power down enable. Read-write. Reset: 0. BIOS: 1. 1=The DCT places the DRAM devices in precharge power down mode when pages are closed as specified by D18F2x248_dct[1:0]_mp[1:0][AggrPDDelay] . 0=The DCT places the DRAM devices in precharge power down mode when pages are closed as specified by D18F2x90_dct[1:0][DynPageCloseEn] .

20	BankSwap: swap bank address. Read-write. Reset: 0. BIOS: See 2.9.5.7. 1=Swap the DRAM bank address bits. IF (D18F2x110[DctSelIntLvEn]==1 && D18F2x110[DctSelIntLvAddr]==100b) THEN normalized address bits (7+n):8 are swapped with normalized address bits used for bank address (See D18F2x80_dct[1:0]) ELSE normalized address bits (8+n):9 are swapped with normalized address bits used for bank address; n is the number of bank address bits for the chip select. For example, if D18F2x110[DctSelIntLvAddr]==101b and D18F2x80_dct[1:0][DimmAddrMap0]==0001b, then normalized address bits 11:9 are swapped with normalized address bits 15:13, when sent to chip select 0. This swap happens before D18F2x94_dct[1:0][BankSwizzleMode] is applied.																				
19:18	Reserved.																				
17:16	MemPhyPllPdMode: memory phy PLL powerdown mode. Read-write. Reset: 00b. BIOS: 10b. These bits control how the memory PLL powers down during self-refresh. If self-refresh is requested for an NB p-state change, then the memory PLL is not powered down. The value of this field should be programmed the same for both DCTs. Memory phy PLL powerdown can only be enabled if NB PLL power down is enabled. See D18F5x128[NbPllPwrDwnRegEn]. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>PLL powerdown is disabled</td> </tr> <tr> <td>01b</td> <td>PLL VCO powerdown during SR</td> </tr> <tr> <td>10b</td> <td>PLL regulator powerdown during SR</td> </tr> <tr> <td>11b</td> <td>PLL VCO and regulator powerdown during SR</td> </tr> </tbody> </table>	Bits	Description	00b	PLL powerdown is disabled	01b	PLL VCO powerdown during SR	10b	PLL regulator powerdown during SR	11b	PLL VCO and regulator powerdown during SR										
Bits	Description																				
00b	PLL powerdown is disabled																				
01b	PLL VCO powerdown during SR																				
10b	PLL regulator powerdown during SR																				
11b	PLL VCO and regulator powerdown during SR																				
15:8	CtrlWordCS[7:0]: control word chip select. Read-write. Reset: 0. Specifies the target DIMM chip selects used for control word programming; Used in conjunction with D18F2x7C_dct[1:0][SendControlWord]. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>02h-00h</td> <td>Reserved</td> </tr> <tr> <td>03h</td> <td>CS0,CS1 is asserted</td> </tr> <tr> <td>0Bh-04h</td> <td>Reserved</td> </tr> <tr> <td>0Ch</td> <td>CS2,CS3 is asserted</td> </tr> <tr> <td>2Fh-0Dh</td> <td>Reserved</td> </tr> <tr> <td>30h</td> <td>CS4,CS5 is asserted</td> </tr> <tr> <td>BFh-31h</td> <td>Reserved</td> </tr> <tr> <td>C0h</td> <td>CS6,CS7 is asserted</td> </tr> <tr> <td>FFh-C1h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	02h-00h	Reserved	03h	CS0,CS1 is asserted	0Bh-04h	Reserved	0Ch	CS2,CS3 is asserted	2Fh-0Dh	Reserved	30h	CS4,CS5 is asserted	BFh-31h	Reserved	C0h	CS6,CS7 is asserted	FFh-C1h	Reserved
Bits	Description																				
02h-00h	Reserved																				
03h	CS0,CS1 is asserted																				
0Bh-04h	Reserved																				
0Ch	CS2,CS3 is asserted																				
2Fh-0Dh	Reserved																				
30h	CS4,CS5 is asserted																				
BFh-31h	Reserved																				
C0h	CS6,CS7 is asserted																				
FFh-C1h	Reserved																				
7:6	Reserved.																				
5	SubMemclkRegDly: Sub-one MEMCLK register delay. Read-write. Reset: 0. 1=The delay through the DIMM register and routing delay is less than 1 MEMCLK. The dram controller does not add 1 MEMCLK to calculate write latency for DDR3 registered DIMMs. 0=The delay through the DIMM register and routing delay is at least 1 MEMCLK. The DRAM controller adds 1 MEMCLK to calculate write latency for DDR3 registered DIMMs.																				
4	Reserved.																				
3	Reserved.																				
2	CSTimingMux67: CS timing mux 6 and 7. Read-write. Reset: 0. BIOS: 2.9.5.4.3. Selects the register used for the setup time of the CS[7:6] pins versus MEMCLK. 0=D18F2x9C_x0000_0004_dct[1:0]_mp[1:0][CsOdtSetup]. 1=D18F2x9C_x0000_0004_dct[1:0]_mp[1:0][AddrCmdSetup].																				
1:0	Reserved.																				

D18F2xAC DRAM Controller Temperature Status

Cold reset: 0000_0000h.

Bits	Description
31:3	Reserved.
2	MemTempHot1: Memory temperature hot, DCT1. See: MemTempHot0.
1	Reserved.
0	MemTempHot0: Memory temperature hot, DCT0. Read; Write-1-to-clear. 1=The EVENT_L pin was asserted indicating the memory temperature exceeded the normal operating limit; the DCT may be throttling the interface to aid in cooling. See D18F2xA4 .

D18F2xF8 P-state Power Information 1

Read-only.

Bits	Description										
31:24	PwrValue3: P3 power value. See PwrValue0. Value: Product-specific.										
23:16	PwrValue2: P2 power value. See PwrValue0. Value: Product-specific.										
15:8	PwrValue1: P1 power value. See PwrValue0. Value: Product-specific.										
7:0	PwrValue0: P0 power value. PwrValue and PwrDiv together specify the expected power draw of a single core in P0 and 1/NumCores of the Northbridge in the NB P-state as specified by MSRC001_00[6B:64][NbPstate] . NumCores is defined to be the number of cores per node at cold reset. Value: Product-specific. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th><u>PwrDiv</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>PwrValue / 1 W, Range: 0 to 255 W</td> </tr> <tr> <td>01b</td> <td>PwrValue / 10 W, Range: 0 to 25.5 W</td> </tr> <tr> <td>10b</td> <td>PwrValue / 100 W, Range: 0 to 2.55 W</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>PwrDiv</u>	<u>Description</u>	00b	PwrValue / 1 W, Range: 0 to 255 W	01b	PwrValue / 10 W, Range: 0 to 25.5 W	10b	PwrValue / 100 W, Range: 0 to 2.55 W	11b	Reserved
<u>PwrDiv</u>	<u>Description</u>										
00b	PwrValue / 1 W, Range: 0 to 255 W										
01b	PwrValue / 10 W, Range: 0 to 25.5 W										
10b	PwrValue / 100 W, Range: 0 to 2.55 W										
11b	Reserved										

D18F2x104 P-state Power Information 3

Read-only.

Bits	Description
31:24	Reserved.
23:16	PwrValue7: P7 power value. See D18F2xF8[PwrValue0] . Value: Product-specific.
15:8	PwrValue6: P6 power value. See D18F2xF8[PwrValue0] . Value: Product-specific.
7:0	PwrValue5: P5 power value. See D18F2xF8[PwrValue0] . Value: Product-specific.

D18F2x10C Swap Interleaved Region Base/Limit

IF ([D18F2x118\[LockDramCfg\]](#)) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0000_0000h. Enables swapping a region below 16G with the same sized region located at the bottom of memory. This register is typically used to map addresses of a graphics frame buffer located below the sub-4GB IO hole to interleaved DRAM in low memory, and is only necessary if the frame buffer normally maps to non-interleaved memory.

- The size of the swapped high region must be a integer multiple of 128M, defined to be {IntLvRgnBaseAddr,

- 000b, 000000h} to {IntLvRgnLmtAddr, 111b, FFFFFFFh}.
- The size of the swapped region must be less than or equal to the alignment of IntLvRgnBaseAddr.
 - E.g. If IntLvRgnBaseAddr=2h then size <=256MB and BIOS programs IntLvRgnLmtAddr <= IntLvRgnBaseAddr + 1h.
 - It is expected that BIOS may program IntLvRgnBaseAddr to a value less than the base address of the graphics frame buffer if realignment is necessary to achieve a larger swap size.
 - The location of the low region is defined to be 0000_0000h to {IntLvRgnLmtAddr - IntLvRgnBaseAddr, 111b, FFFFFFFh}.
 - The swapped region must be all DRAM. I.e. No IO hole.
 - Channel interleaving must be enabled and the DCTs must be of unequal size.
 - Swapping must not be enabled on more than one node, and D18F1x[144:140,44:40][DramBase] must be zero.
 - See D18F2x110[DctSelIntLvEn]. See 2.9.7 [Memory Hoisting].

Bits	Description
31:27	Reserved.
26:20	IntLvRgnSize[33:27]: Interleave swap region size bits[33:27]. Interleave swap region size [33:27].
19:18	Reserved.
17:11	IntLvRgnLmtAddr[33:27]: Interleave swap region limit address bits[33:27]. Interleave swap region limit address [33:27].
10	Reserved.
9:3	IntLvRgnBaseAddr[33:27]: Interleave swap region base address bits[33:27]. Interleave swap region base address [33:27].
2:1	Reserved.
0	IntLvRgnSwapEn: Interleave region swap enable. 1=Enables swapping a region from the top of memory to the bottom of DRAM space.

D18F2x110 DRAM Controller Select Low

Reset: 0000_0000h.

Bits	Description
31:11	DctSelBaseAddr[47:27]: DRAM controller select base address bits[47:27]. IF (D18F2x118[Lock-DramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Delineates the address range of the two DCTs by specifying the base address of the upper address range. See DctSelHiRngEn.
10	MemCleared: memory cleared. Read-only. 1=Memory has been cleared since the last warm reset. This bit is set by MemClrInit. See MemClrInit.
9	MemClrBusy: memory clear busy. Read-only. 1=The memory clear operation in either of the DCTs is in progress. Reads or writes to DRAM while the memory clear operation is in progress result in undefined behavior.
8	DramEnable: DRAM enabled. Read-only. 1=All of the used DCTs are initialized (see 2.9.5.8 [DRAM Device and Controller Initialization]) or have exited from self refresh (D18F2x90_dct[1:0][ExitSelfRef] transitions from 1 to 0).

7:6	<p>DctSelIntLvAddr[1:0]: DRAM controller select channel interleave address bit. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. IF (DctSelIntLvEn) THEN BIOS: 00b. ENDIF. Specifies how interleaving is selected between the DCTs. In all cases, if the select function is low then DCT0 is selected; if the select function is high then DCT1 is selected. DctSelIntLvAddr[2:0] = {D18F2x114[DctSelIntLvAddr[2], D18F2x110[DctSelIntLvAddr[1:0]]}. The select functions are defined as follows:</p> <table border="1" data-bbox="280 401 1055 709"> <thead> <tr> <th>DctSelIntLvAddr</th> <th>DCT Select Function</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Address bit 6.</td> </tr> <tr> <td>001b</td> <td>Address bit 12.</td> </tr> <tr> <td>010b</td> <td>Hash: exclusive OR of address bits[20:16, 6].</td> </tr> <tr> <td>011b</td> <td>Hash: exclusive OR of address bits[20:16, 9].</td> </tr> <tr> <td>100b</td> <td>Address bit 8.</td> </tr> <tr> <td>101b</td> <td>Address bit 9.</td> </tr> <tr> <td>110b</td> <td>Reserved.</td> </tr> <tr> <td>111b</td> <td>Reserved.</td> </tr> </tbody> </table> <p>If the internal GPU is enabled, only encodings 001b, 100b, and 101b are supported.</p>	DctSelIntLvAddr	DCT Select Function	000b	Address bit 6.	001b	Address bit 12.	010b	Hash: exclusive OR of address bits[20:16, 6].	011b	Hash: exclusive OR of address bits[20:16, 9].	100b	Address bit 8.	101b	Address bit 9.	110b	Reserved.	111b	Reserved.
DctSelIntLvAddr	DCT Select Function																		
000b	Address bit 6.																		
001b	Address bit 12.																		
010b	Hash: exclusive OR of address bits[20:16, 6].																		
011b	Hash: exclusive OR of address bits[20:16, 9].																		
100b	Address bit 8.																		
101b	Address bit 9.																		
110b	Reserved.																		
111b	Reserved.																		
5	<p>DctDatIntLv: DRAM controller data interleave enable. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. BIOS: Set if ECC is enabled. 1=DRAM data bits from every two consecutive 64-bit DRAM lines are interleaved in the ECC calculation such that a dead bit of a DRAM device is correctable.</p>																		
4	Reserved.																		
3	<p>MemClrInit: memory clear initialization. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. 1=The node writes 0's to all locations of system memory attached to the node and sets the MemCleared bit. The status of the memory clear operation can be determined by reading the MemClrBusy and MemCleared bits. This command is ignored if MemClrBusy=1 when the command is received. DramEnable must be set before setting MemClrInit. The memory prefetcher must be disabled by setting D18F2x11C[PrefIoDis] and D18F2x11C[PrefCpuDis] before memory clear initialization and then can be re-enabled when MemCleared=1.</p>																		
2	<p>DctSelIntLvEn: DRAM controller interleave enable. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. BIOS: See 2.9.5.7. 1=Channel interleave is enabled; DctSelIntLvAddr specifies which address bit is used to select between DCT0 and DCT1; this applies from the base system memory address of the node (specified by D18F1x[144:140,44:40] [DRAM Base/Limit]) to DctSelBaseAddr (if enabled). If the amount of memory connected to each of the DCTs is different, then channel interleaving may be supported across the address range that includes both DCTs, the top of which is specified by DctSelBaseAddr; the remainder of the address space, above DctSelBaseAddr, would then be allocated to only the DCT connected to the larger amount of memory, specified by DctSelHi.</p>																		
1	<p>DctSelHi: DRAM controller high select. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. If DctSelHiRngEn is set, this specifies which DCT receives accesses with addresses in the high range (greater than or equal to DctSelBaseAddr). 0=High addresses go to DCT0. 1=High addresses go to DCT1.</p>																		
0	<p>DctSelHiRngEn: DRAM controller select high range enable. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. 1=Enables addresses greater than or equal to DctSelBaseAddr[47:27] to be used to select between DCT0 and DCT1; DctSelHi specifies which DCT occupies the high range.</p>																		

D18F2x114 DRAM Controller Select High

IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0000_0000h.

Bits	Description
31:10	DctSelBaseOffset[47:26]: DRAM controller select base offset address bits[47:26]. When D18F2x110[DctSelHiRngEn]=1, this value is subtracted from the physical address of certain transactions before being passed to the DCT. See 2.9.7.2 [DctSelBaseOffset Programming].
9	DctSelIntLvAddr[2]: DRAM controller select channel interleave address bit. IF (D18F2x110[DctSelIntLvEn]) THEN BIOS: 1. ENDIF. See D18F2x110[DctSelIntLvAddr[1:0]].
8:0	Reserved.

D18F2x118 Memory Controller Configuration Low

Fields in this register (bits[17:0]) indicate priority of request types. Variable priority requests enter the memory controller as medium priority and are promoted to high priority if they have not been serviced in the time specified by MctVarPriCntLmt. This feature may be useful for isochronous IO traffic. If isochronous traffic is specified to be high priority, it may have an adverse effect on the bandwidth and performance of the devices associated with the other types of traffic. However, if isochronous traffic is specified as medium priority, the processor may not meet the isochronous bandwidth and latency requirements. The variable priority allows the memory controller to optimize DRAM transactions until isochronous traffic reaches a time threshold and must be serviced more quickly.

Bits	Description																																				
31:28	MctVarPriCntLmt: variable priority time limit. Read-write. Reset: 0000b. BIOS: 0000b. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>80ns</td> <td>1000b</td> <td>720ns</td> </tr> <tr> <td>0001b</td> <td>160ns</td> <td>1001b</td> <td>800ns</td> </tr> <tr> <td>0010b</td> <td>240ns</td> <td>1010b</td> <td>880ns</td> </tr> <tr> <td>0011b</td> <td>320ns</td> <td>1011b</td> <td>960ns</td> </tr> <tr> <td>0100b</td> <td>400ns</td> <td>1100b</td> <td>1040ns</td> </tr> <tr> <td>0101b</td> <td>480ns</td> <td>1101b</td> <td>1120ns</td> </tr> <tr> <td>0110b</td> <td>560ns</td> <td>1110b</td> <td>1200ns</td> </tr> <tr> <td>0111b</td> <td>640ns</td> <td>1111b</td> <td>1280ns</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	0000b	80ns	1000b	720ns	0001b	160ns	1001b	800ns	0010b	240ns	1010b	880ns	0011b	320ns	1011b	960ns	0100b	400ns	1100b	1040ns	0101b	480ns	1101b	1120ns	0110b	560ns	1110b	1200ns	0111b	640ns	1111b	1280ns
Bits	Description	Bits	Description																																		
0000b	80ns	1000b	720ns																																		
0001b	160ns	1001b	800ns																																		
0010b	240ns	1010b	880ns																																		
0011b	320ns	1011b	960ns																																		
0100b	400ns	1100b	1040ns																																		
0101b	480ns	1101b	1120ns																																		
0110b	560ns	1110b	1200ns																																		
0111b	640ns	1111b	1280ns																																		
27	MctEccDisLatOptEn: memory controller ecc disabled latency optimization. Read-write. Reset: 0. BIOS: 1. 1=Enable bypass of memory ECC checking logic. BIOS must program this bit before any memory accesses are issued from the processor.																																				
26:24	McqHiPriByPassMax: memory controller high priority bypass max. Read-write. Reset: 100b. Specifies the number of times a medium- or low-priority DRAM request may be bypassed by high-priority DRAM requests.																																				
23	Reserved.																																				
22:20	McqMedPriByPassMax: memory controller medium bypass low priority max. Read-write. Reset: 100b. Specifies the number of times a low-priority DRAM request may be bypassed by medium-priority DRAM requests.																																				

19	<p>LockDramCfg. Read; write-1-only. Reset: 0. BIOS: See 2.9.8 [DRAM CC6/PC6 Storage], 2.5.3.2.3.3 [Core C6 (CC6) State].</p> <p>The following registers are read-only if LockDramCfg=1; otherwise the access type is specified by the register:</p> <ul style="list-style-type: none"> • D18F1xF0 [DRAM Hole Address] • D18F2x[5C:40]_dct[1:0] [DRAM CS Base Address] • D18F2x[6C:60]_dct[1:0] [DRAM CS Mask] • D18F2x80_dct[1:0] [DRAM Bank Address Mapping] • D18F2x10C [Swap Interleaved Region Base/Limit] • D18F2x110 [DRAM Controller Select Low] • D18F2x114 [DRAM Controller Select High] • D18F4x128[CoreStateSaveDestNode] • D18F1x[144:140,44:40] [DRAM Base/Limit] • D18F1x120 [DRAM Base System Address] • D18F1x124 [DRAM Limit System Address] • D18F2x118[CC6SaveEn] 										
18	<p>CC6SaveEn. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0. 1=CC6 save area is enabled. See 2.5.3.2.7 [BIOS Requirements for Initialization]. BIOS: (D18F4x118[PwrGateEnCstAct0] D18F4x118[PwrGateEnCstAct1] D18F4x11C[PwrGateEnCstAct2]).</p>										
17:16	<p>MctPriScrub: scrubber priority. Read-write. Reset: 00b.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Medium</td> </tr> <tr> <td>01b</td> <td>Reserved</td> </tr> <tr> <td>10b</td> <td>High</td> </tr> <tr> <td>11b</td> <td>Variable</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Medium	01b	Reserved	10b	High	11b	Variable
<u>Bits</u>	<u>Description</u>										
00b	Medium										
01b	Reserved										
10b	High										
11b	Variable										
15:14	<p>MctPriTrace: trace-mode request priority. See: MctPriCpuRd. Read-write. Reset: 10b. This must be set to high.</p>										
13:12	<p>MctPriIsoc: display refresh read priority. See: MctPriCpuRd. Read-write. Reset: 10b.</p>										
11:10	<p>MctPriWr: default write priority. See: MctPriCpuRd. Read-write. Reset: 01b.</p>										
9:8	<p>MctPriDefault: default non-write priority. See: MctPriCpuRd. Read-write. Reset: 00b.</p>										
7:6	<p>MctPriIsocWr: IO write with the isoch bit set priority. See: MctPriCpuRd. Read-write. Reset: 00b. This does not apply to isochronous traffic that is classified as display refresh.</p>										
5:4	<p>MctPriIsocRd: IO read with the isoch bit set priority. See: MctPriCpuRd. Read-write. Reset: 10b. This does not apply to isochronous traffic that is classified as display refresh.</p>										
3:2	<p>MctPriCpuWr: CPU write priority. See: MctPriCpuRd. Read-write. Reset: 01b.</p>										
1:0	<p>MctPriCpuRd: CPU read priority. Read-write. Reset: 00b.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Medium</td> </tr> <tr> <td>01b</td> <td>Low</td> </tr> <tr> <td>10b</td> <td>High</td> </tr> <tr> <td>11b</td> <td>Variable</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Medium	01b	Low	10b	High	11b	Variable
<u>Bits</u>	<u>Description</u>										
00b	Medium										
01b	Low										
10b	High										
11b	Variable										

D18F2x11C Memory Controller Configuration High

The two main functions of this register are to control write bursting and memory prefetching.

Write bursting. DctWrLimit and MctWrLimit specify how writes may be burst from the MCT into the DCT to improve DRAM efficiency. When the number of writes in the MCT reaches the value specified in MctWrLimit, then they are all burst to the DCTs at once. Prior to reaching the watermark, a limited number of writes can be passed to the DCTs (specified by DctWrLimit), tagged as low priority, for the DCTs to complete when otherwise idle. Rules regarding write bursting:

- Write bursting mode only applies to low-priority writes. Medium and high priority writes are not withheld from the DCTs for write bursting.
- If write bursting is enabled, writes stay in the MCQ until the threshold specified by MctWrLimit is reached.
- Once the threshold is reached, all writes in MCQ are converted to medium priority.
- Any write in MCQ that matches the address of a subsequent access is promoted to either medium priority or the priority of the subsequent access, whichever is higher.
- DctWrLimit only applies to low-priority writes.

Memory prefetching. The MCT prefetcher detects stride patterns in the stream of requests and then, for predictable stride patterns, generates prefetch requests. A stride pattern is a pattern of requests through system memory that are the same number of cachelines apart. The prefetcher supports strides of -4 to +4 cachelines, which can include alternating patterns (e.g. +1, +2, +1, +2), and can prefetch 1, 2, 3, 4, or 5 cachelines ahead, depending on the confidence. In addition, a fixed stride mode (non-alternating) may be used for IO requests which often have fixed stride patterns. This mode bypasses the stride predictor such that CPU-access stride predictions are not adversely affected by IO streams.

The MCT tracks several stride patterns simultaneously. Each of these has a confidence level associated with it that varies as follows:

- Each time a request is received that matches the stride pattern, the confidence level increases by one.
- Each time a request is received within +/- 4 cachelines of the last requested cacheline in the pattern that does not match the pattern, then the confidence level decreases by one.
- When the confidence level reaches the saturation point specified by PrefConfSat, then it no-longer increments.

Each request that is not within +/- 4 cachelines of the last requested cacheline line of all the stride patterns tracked initiates a new stride pattern by displacing one of the existing least-recently-used stride patterns.

The memory prefetcher uses an adaptive prefetch scheme to adjust the prefetch distance based upon the buffer space available for prefetch request data. The adaptive scheme counts the total number of prefetch requests and the number of prefetch requests that cannot return data because of buffer availability. After every 16 prefetch requests, the prefetcher uses the following rules to adjust the prefetch distance:

- If the ratio of prefetch requests that cannot return data to total prefetch requests is greater than or equal to `D18F2x1B0[AdapPrefMissRatio]` then the prefetch distance is reduced by `D18F2x1B0[AdapPrefNegativeStep]`.
- If the ratio of prefetch requests that cannot return data to total prefetch requests is less than `D18F2x1B0[AdapPrefMissRatio]` then the prefetch distance is increased by `D18F2x1B0[AdapPrefPositiveStep]`.
- If the adjusted prefetch distance is greater than the prefetch distance defined for the current confidence level, the prefetch distance for the current confidence level is used.

The adaptive prefetch scheme supports fractional prefetch distances by alternating between two whole number prefetch distances. For example a prefetch distance of 1.25 causes a prefetch distance sequence of: 1, 1, 1, 2, 1, 1, 1, 2.

Bits	Description
31	Reserved.

30	FlushWr: flush writes command. Read; write-1-only; cleared-by-hardware. Reset: 0. Setting this bit causes write bursting to be canceled and all outstanding writes to be flushed to DRAM. This bit is cleared when all writes are flushed to DRAM.										
29	FlushWrOnStpGnt: flush writes on stop-grant. Read-write. Reset: 0. BIOS: ~D18F2x1B4[FlushWrOnS3StpGnt]. 1=Causes write bursting to be canceled and all outstanding writes to be flushed to DRAM when in the stop-grant state.										
28	PrefDramTrainMode: prefetch DRAM training mode. Read-write; cleared-by-hardware. Reset: 0. 1=Enable DRAM training mode. Hardware clears this bit when the prefetch request limit is reached. Writing a zero to this bit clears the prefetch buffer and disables the DRAM training mode prefetcher. BIOS must write a zero to this bit after training is complete. See 2.9.5.9.6 [Continuous Pattern Generation].										
27:25	PrefThreeConf: prefetch three-ahead confidence. Read-write. Reset: 100b. BIOS: 110b. Confidence level required in order to prefetch three cachelines ahead (same encoding as PrefTwoConf below).										
24:22	PrefTwoConf: prefetch two-ahead confidence. Read-write. Reset: 011b. BIOS: 011b. Confidence level required in order to prefetch two cachelines ahead. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0</td> </tr> <tr> <td>110b-001b</td> <td><PrefTwoConf*2></td> </tr> <tr> <td>111b</td> <td>14</td> </tr> </tbody> </table>	Bits	Description	000b	0	110b-001b	<PrefTwoConf*2>	111b	14		
Bits	Description										
000b	0										
110b-001b	<PrefTwoConf*2>										
111b	14										
21:20	PrefOneConf: prefetch one-ahead confidence. Read-write. Reset: 10b. BIOS: 10b. Confidence level required in order to prefetch one ahead (0 through 3).										
19:18	PrefConfSat: prefetch confidence saturation. Read-write. Reset: 00b. BIOS: 00b. Specifies the point at which prefetch confidence level saturates and stops incrementing. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>15</td> </tr> <tr> <td>01b</td> <td>7</td> </tr> <tr> <td>10b</td> <td>3</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	15	01b	7	10b	3	11b	Reserved
Bits	Description										
00b	15										
01b	7										
10b	3										
11b	Reserved										
17:16	PrefFixDist: prefetch fixed stride distance. Read-write. Reset: 00b. Specifies the distance to prefetch ahead if in fixed stride mode. 00b=1 cacheline; 01b=2 cachelines; 10b=3 cachelines; 11b=4 cachelines.										
15	PrefFixStrideEn: prefetch fixed stride enable. Read-write. Reset: 0. 1=The prefetch stride for all requests (CPU and IO) is fixed (non-alternating).										
14	PrefIoFixStrideEn: Prefetch IO fixed stride enable. Read-write. Reset: 0. 1=The prefetch stride for IO requests is fixed (non-alternating).										
13	PrefIoDis: prefetch IO-access disable. Read-write. Reset: 1. BIOS: 0. 1=Disables IO requests from triggering prefetch requests.										
12	PrefCpuDis: prefetch CPU-access disable. Read-write. Reset: 1. BIOS: 0. 1=Disables CPU requests from triggering prefetch requests.										
11:7	MctPrefReqLimit: memory controller prefetch request limit. Read-write. Reset: 1Eh. Specifies the maximum number of outstanding prefetch requests allowed. See D18F3x78 for restrictions on this field.										

6:2	<p>MctWrLimit: memory controller write-burst limit. Read-write. Reset: 1Fh. BIOS: 0Ch. Specifies the number of writes in the memory controller queue before they are burst into the DCTs.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>16</td> </tr> <tr> <td>0Fh-01h</td> <td><16-MctWrLimit></td> </tr> <tr> <td>1Fh-10h</td> <td>Write bursting disabled</td> </tr> </tbody> </table>	Bits	Description	00h	16	0Fh-01h	<16-MctWrLimit>	1Fh-10h	Write bursting disabled		
Bits	Description										
00h	16										
0Fh-01h	<16-MctWrLimit>										
1Fh-10h	Write bursting disabled										
1:0	<p>DctWrLimit: DRAM controller write limit. Read-write. Reset: 00b. BIOS: 01b. Specifies the maximum number of writes allowed in the DCT queue when write bursting is enabled, prior to when the number of writes in MCQ exceeds the watermark specified by MctWrLimit.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>0</td> </tr> <tr> <td>01b</td> <td>2</td> </tr> <tr> <td>10b</td> <td>4</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	0	01b	2	10b	4	11b	Reserved
Bits	Description										
00b	0										
01b	2										
10b	4										
11b	Reserved										

D18F2x1B0 Extended Memory Controller Configuration Low

The main function of this register is to control the memory prefetcher. See [D18F2x11C \[Memory Controller Configuration High\]](#) about the adaptive prefetch scheme.

Table 176: D18F2x1B[4:0] Recommended Settings

Condition	D18F2x1B0	D18F2x1B4	
	DcqBwThrotWm	DcqBwThrotWm1	DcqBwThrotWm2
DdrRate	0h	3h	4h
667		3h	5h
800		4h	6h
1066		5h	8h
1333		6h	9h
1600		7h	Ah
1866			

Bits	Description								
31:28	<p>DcqBwThrotWm: dcq bandwidth throttle watermark. Read-write. Reset: 3h. BIOS: Table 176. Specifies the number of outstanding DRAM read requests before new DRAM prefetch requests and speculative prefetch requests are throttled. 0h=Throttling is determined by D18F2x1B4[DcqBwThrotWm1, DcqBwThrotWm2]. Legal values are 0h through Ch.</p>								
27:25	<p>PrefFiveConf: prefetch five-ahead confidence. Read-write. Reset: 110b. BIOS: 111b. Confidence level required in order to prefetch five cachelines ahead.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0</td> </tr> <tr> <td>110b-001b</td> <td><PrefFiveConf*2></td> </tr> <tr> <td>111b</td> <td>14</td> </tr> </tbody> </table>	Bits	Description	000b	0	110b-001b	<PrefFiveConf*2>	111b	14
Bits	Description								
000b	0								
110b-001b	<PrefFiveConf*2>								
111b	14								

24:22	<p>PrefFourConf: prefetch four-ahead confidence. Read-write. Reset: 101b. BIOS: 111b. Confidence level required in order to prefetch four cachelines ahead.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0</td> </tr> <tr> <td>110b-001b</td> <td><PrefFourConf*2></td> </tr> <tr> <td>111b</td> <td>14</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	000b	0	110b-001b	<PrefFourConf*2>	111b	14				
<u>Bits</u>	<u>Description</u>												
000b	0												
110b-001b	<PrefFourConf*2>												
111b	14												
21	Reserved.												
20	<p>DbIPrefEn: double prefetch enable. Read-write. Reset: 0. 1=The memory prefetcher only generates prefetch requests when it is able to generate a pair of prefetch requests to consecutive cache lines. BIOS must clear D18F2x11C[PrefDramTrainMode] prior to setting this bit.</p>												
19:18	Reserved. Reset: 00b												
17:13	Reserved. Reset: 11100b.												
12	<p>EnSplitDctLimits: split DCT write limits enable. Read-write. Reset: 0. BIOS: 1. 1=The number of writes specified by D18F2x11C[DctWrLmt and MctWrLmt] is per DCT. 0=The number of writes specified by D18F2x11C[DctWrLmt and MctWrLmt] is total writes independent of DCT. Setting this bit also affects the encoding of D18F2x11C[DctWrLmt].</p>												
11	<p>DisIoCohPref: disable coherent prefetched for IO. Read-write. Reset: 0. 1=Probes are not generated for prefetches generated for reads from IO devices.</p>												
10:8	<p>CohPrefPrbLmt: coherent prefetch probe limit. Read-write. Reset: 000b. BIOS: 000b. Specifies the maximum number of probes that can be outstanding for memory prefetch requests.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Probing disabled for memory prefetch requests</td> </tr> <tr> <td>001b</td> <td>4 outstanding probes</td> </tr> <tr> <td>010b</td> <td>8 outstanding probes</td> </tr> <tr> <td>011b</td> <td>Reserved.</td> </tr> <tr> <td>1xxb</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	000b	Probing disabled for memory prefetch requests	001b	4 outstanding probes	010b	8 outstanding probes	011b	Reserved.	1xxb	Reserved
<u>Bits</u>	<u>Description</u>												
000b	Probing disabled for memory prefetch requests												
001b	4 outstanding probes												
010b	8 outstanding probes												
011b	Reserved.												
1xxb	Reserved												
7:6	Reserved.												
5:4	<p>AdapPrefNegativeStep: adaptive prefetch negative step. Read-write. Reset: 00b. BIOS: 00b. Specifies the step size that the adaptive prefetch scheme uses when decreasing the prefetch distance.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>2/16</td> </tr> <tr> <td>01b</td> <td>4/16</td> </tr> <tr> <td>10b</td> <td>8/16</td> </tr> <tr> <td>11b</td> <td>16/16</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	2/16	01b	4/16	10b	8/16	11b	16/16		
<u>Bits</u>	<u>Description</u>												
00b	2/16												
01b	4/16												
10b	8/16												
11b	16/16												

3:2	AdapPrefPositiveStep: adaptive prefetch positive step. Read-write. Reset: 00b. BIOS: 00b. Specifies the step size that the adaptive prefetch scheme uses when increasing the prefetch distance.										
	<table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1/16</td> </tr> <tr> <td>01b</td> <td>2/16</td> </tr> <tr> <td>10b</td> <td>4/16</td> </tr> <tr> <td>11b</td> <td>8/16</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	1/16	01b	2/16	10b	4/16	11b	8/16
<u>Bits</u>	<u>Description</u>										
00b	1/16										
01b	2/16										
10b	4/16										
11b	8/16										
1:0	AdapPrefMissRatio: adaptive prefetch miss ratio. Read-write. Reset: 00b. BIOS: 01b. Specifies the ratio of prefetch requests that do not have data buffer available to the total number of prefetch requests at which the adaptive prefetch scheme begins decreasing the prefetch distance.										
	<table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1/16</td> </tr> <tr> <td>01b</td> <td>2/16</td> </tr> <tr> <td>10b</td> <td>4/16</td> </tr> <tr> <td>11b</td> <td>8/16</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	1/16	01b	2/16	10b	4/16	11b	8/16
<u>Bits</u>	<u>Description</u>										
00b	1/16										
01b	2/16										
10b	4/16										
11b	8/16										

D18F2x1B4 Extended Memory Controller Configuration High Register

Bits	Description
31	FlushOnMmioWrEn: flush on mmio write enable. Read-write. Reset: 0. 1=Any CPU-sourced MMIO write that matches D18F1x[1CC:180,BC:80] causes the memory controller data buffers to be flushed to memory.
30:28	S3SmafId: S3 SMAF id. Read-write. Reset: 100b. SMAF encoding of D18F3x[84:80] corresponding to the ACPI S3 state when FlushWrOnS3StpGnt=1. Reserved when FlushWrOnS3StpGnt=0.
27	FlushWrOnS3StpGnt: flush write on S3 stop grant. Read-write. Reset: 0. BIOS: 1. 1=Write bursting is canceled and all outstanding writes are flushed to DRAM when in the stop-grant state and the SMAF code is equal to S3SmafId, indicating entry into the ACPI S3 state. See D18F2xA8_dct[1:0][FastSelfRefEntryDis] , D18F2x11C[FlushWrOnStpGnt] .
26	EnSplitMctDatBuffers: enable split MCT data buffers. Read-write. Reset: 0. BIOS: 1. 1=Enable resource allocation into the split buffer resources BIOS must program this bit before any memory accesses are issued from the processor.
25:23	Reserved.
22	SpecPrefDisWm1: speculative prefetch disable watermark 1. Read-write. Reset: 0. 0=Disable speculative prefetches at the DcqBwThrotWm2 limit. 1=Disable speculative prefetches at the DcqBwThrotWm1 limit. See also D18F2x1B0[SpecPrefDis] .
21	RegionAlloWm2: region prefetch allocate watermark 2. Read-write. Reset: 0. See DemandAlloWm2.
20	RegionPropWm2: region prefetch propagate watermark 2. Read-write. Reset: 0. See DemandPropWm2.
19	StrideAlloWm2: stride prefetch allocate watermark 2. Read-write. Reset: 1. See DemandAlloWm2.
18	StridePropWm2: stride prefetch propagate watermark 2. Read-write. Reset: 1. See DemandPropWm2.

17	DemandAlloWm2: demand request allocate watermark 2. Read-write. Reset: 1. Specifies the behavior from the DcqBwThrotWm1 limit to the DcqBwThrotWm2 limit. 0=Requests do not allocate a new entry. 1=Requests allocate a new entry; defined only if (DemandAlloWm1 & DemandPropWm2).
16	DemandPropWm2: demand request propagate watermark 2. Read-write. Reset: 1. Specifies the behavior from the DcqBwThrotWm1 limit to the DcqBwThrotWm2 limit. 0=Requests do not update existing entries. 1=Requests update existing entries; defined only if (DemandPropWm1=1).
15	RegionAlloWm1: region prefetch allocate watermark 1. Read-write. Reset: 0. See DemandAlloWm1.
14	RegionPropWm1: region prefetch propagate watermark 1. Read-write. Reset: 1. See DemandPropWm1.
13	StrideAlloWm1: stride prefetch allocate watermark 1. Read-write. Reset: 1. See DemandAlloWm1.
12	StridePropWm1: stride prefetch propagate watermark 1. Read-write. Reset: 1. See DemandPropWm1.
11	DemandAlloWm1: demand request allocate watermark 1. Read-write. Reset: 1. Specifies the behavior prior to the DcqBwThrotWm1 limit. 0=Requests do not allocate a new entry. 1=Requests allocate a new entry; defined only if (DemandPropWm1=1).
10	DemandPropWm1: demand request propagate watermark 1. Read-write. Reset: 1. Specifies the behavior prior to the DcqBwThrotWm1 limit. 0=Requests do not update existing entries. 1=Requests update existing entries.
9:5	DcqBwThrotWm2: DCQ bandwidth throttle watermark 2. Read-write. Reset: 06h. BIOS: Table 176 . Specifies a prefetch throttling watermark based on the number of outstanding DRAM read requests. This field is reserved when D18F2x1B0[DcqBwThrotWm] != 0. When throttling is enabled, if the number of outstanding DRAM read requests exceeds DcqBwThrotWm2 both request allocate and propagate are blocked and new prefetches are disabled. When throttling is enabled, DcqBwThrotWm2 should be programmed to a value greater than DcqBwThrotWm1. 0h=Throttling is disabled. Legal values are 0h through 18h.
4:0	DcqBwThrotWm1: DCQ bandwidth throttle watermark 1. Read-write. Reset: 03h. BIOS: Table 176 . Specifies a prefetch throttling watermark based on the number of outstanding DRAM read requests. This field is reserved when D18F2x1B0[DcqBwThrotWm] != 0. 0h=Throttling is disabled. Legal values are 0h through 18h.

D18F2x200_dct[1:0]_mp[1:0] DRAM Timing 0

Reset: 0F05_0505h. See [2.9.1 \[DCT Configuration Registers\]](#).

Bits	Description								
31:30	Reserved.								
29:24	Tras: row active strobe. Read-write. Specifies the minimum time in memory clock cycles from an activate command to a precharge command, both to the same chip select bank.								
	<table border="0"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>07h-00h</td> <td>Reserved</td> </tr> <tr> <td>2Ah-08h</td> <td><Tras> clocks</td> </tr> <tr> <td>3Fh-2Bh</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	07h-00h	Reserved	2Ah-08h	<Tras> clocks	3Fh-2Bh	Reserved
<u>Bits</u>	<u>Description</u>								
07h-00h	Reserved								
2Ah-08h	<Tras> clocks								
3Fh-2Bh	Reserved								

23:21	Reserved.								
20:16	<p>Trp: row precharge time. Read-write. Specifies the minimum time in memory clock cycles from a precharge command to an activate command or auto refresh command, both to the same bank.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>04h-00h</td> <td>Reserved</td> </tr> <tr> <td>13h-05h</td> <td><Trp> clocks</td> </tr> <tr> <td>1Fh-14h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	04h-00h	Reserved	13h-05h	<Trp> clocks	1Fh-14h	Reserved
<u>Bits</u>	<u>Description</u>								
04h-00h	Reserved								
13h-05h	<Trp> clocks								
1Fh-14h	Reserved								
15:13	Reserved.								
12:8	<p>Trcd: RAS to CAS delay. Read-write. Specifies the time in memory clock cycles from an activate command to a read/write command, both to the same bank.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>01h-00h</td> <td>Reserved</td> </tr> <tr> <td>13h-02h</td> <td><Trcd> clocks</td> </tr> <tr> <td>1Fh-14h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	01h-00h	Reserved	13h-02h	<Trcd> clocks	1Fh-14h	Reserved
<u>Bits</u>	<u>Description</u>								
01h-00h	Reserved								
13h-02h	<Trcd> clocks								
1Fh-14h	Reserved								
7:5	Reserved.								
4:0	<p>Tcl: CAS latency. Read-write. Specifies the time in memory clock cycles from the CAS assertion for a read cycle until data return (from the perspective of the DRAM devices).</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>04h-00h</td> <td>Reserved</td> </tr> <tr> <td>10h-05h</td> <td><Tcl> clocks</td> </tr> <tr> <td>1Fh-11h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	04h-00h	Reserved	10h-05h	<Tcl> clocks	1Fh-11h	Reserved
<u>Bits</u>	<u>Description</u>								
04h-00h	Reserved								
10h-05h	<Tcl> clocks								
1Fh-11h	Reserved								

D18F2x204_dct[1:0]_mp[1:0] DRAM Timing 1

Reset: 0400_040Bh. See 2.9.1 [DCT Configuration Registers].

Bits	Description										
31:28	Reserved.										
27:24	<p>Trtp: read CAS to precharge time. Read-write. Specifies the earliest time in memory clock cycles a page can be closed after having been read. Satisfying this parameter ensures read data is not lost due to a premature precharge.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>3h-0h</td> <td>Reserved</td> </tr> <tr> <td>Bh-4h</td> <td><Trtp> clocks</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	3h-0h	Reserved	Bh-4h	<Trtp> clocks	Fh-Ch	Reserved		
<u>Bits</u>	<u>Description</u>										
3h-0h	Reserved										
Bh-4h	<Trtp> clocks										
Fh-Ch	Reserved										
23:22	Reserved.										
21:16	<p>FourActWindow: four bank activate window. Read-write. Specifies the rolling tFAW window in memory clock cycles during which no more than 4 banks in an 8-bank device are activated.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>No tFAW window restriction</td> </tr> <tr> <td>05h-01h</td> <td>Reserved</td> </tr> <tr> <td>2Ch-06h</td> <td><FourActWindow> clocks</td> </tr> <tr> <td>3Fh-2Dh</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00h	No tFAW window restriction	05h-01h	Reserved	2Ch-06h	<FourActWindow> clocks	3Fh-2Dh	Reserved
<u>Bits</u>	<u>Description</u>										
00h	No tFAW window restriction										
05h-01h	Reserved										
2Ch-06h	<FourActWindow> clocks										
3Fh-2Dh	Reserved										
15:12	Reserved.										

11:8	Trrd: row to row delay (or RAS to RAS delay). Read-write. Specifies the minimum time in memory clock cycles between activate commands to different chip select banks. <u>Bits</u> 0h 9h-1h Fh-Ah	<u>Description</u> Reserved <Trrd> clocks Reserved
7:6	Reserved.	
5:0	Trc: row cycle time. Read-write. Specifies the minimum time in memory clock cycles from and activate command to another activate command or an auto refresh command, all to the same chip select bank. <u>Bits</u> 09h-00h 3Ah-0Ah 3Fh-3Bh	<u>Description</u> Reserved <Trc> clocks Reserved

D18F2x208_dct[1:0] DRAM Timing 2

See 2.9.1 [DCT Configuration Registers].

Bits	Description	
31:27	Reserved.	
26:24	Trfc3: auto refresh row cycle time for CS 6 and 7. See: Trfc0.	
23:19	Reserved.	
18:16	Trfc2: auto refresh row cycle time for CS 4 and 5. See: Trfc0.	
15:11	Reserved.	
10:8	Trfc1: auto refresh row cycle time for CS 2 and 3. See: Trfc0.	
7:3	Reserved.	
2:0	Trfc0: auto refresh row cycle time for CS 0 and 1. Read-write. Reset: 100b. Specifies the minimum time from a refresh command to the next valid command, except NOP or DES. The recommended programming of this register varies based on DRAM density and speed. <u>Bits</u> 000b 001b 010b 011b 100b 101b 111b-110b	<u>Description</u> Reserved 90 ns (all speeds, 512 Mbit) 110 ns (all speeds, 1 Gbit) 160 ns (all speeds, 2 Gbit) 300 ns (all speeds, 4 Gbit) 350 ns (all speeds, 8 Gbit) Reserved

D18F2x20C_dct[1:0]_mp[1:0] DRAM Timing 3

See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:18	Reserved.

17:16	WrDqDqsEarly: DQ and DQS write early. Read-Write. Reset: 0. Specifies the DQ and DQS launch timing for write commands relative to the Tcwl MEMCLK. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>0 MEMCLK early</td> </tr> <tr> <td>01b</td> <td>0.5 MEMCLK early</td> </tr> <tr> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	0 MEMCLK early	01b	0.5 MEMCLK early	10b	Reserved	11b	Reserved
Bits	Description										
00b	0 MEMCLK early										
01b	0.5 MEMCLK early										
10b	Reserved										
11b	Reserved										
15:12	Reserved.										
11:8	Twtr: internal DRAM write to read command delay. Read-write. Reset: 4h. Specifies the minimum number of memory clock cycles from a write operation to a read operation, both to the same chip select. This is measured from the rising clock edge following last non-masked data strobe of the write to the rising clock edge of the next read command. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>3h-0h</td> <td>Reserved</td> </tr> <tr> <td>Bh-4h</td> <td><Twtr> clocks</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	3h-0h	Reserved	Bh-4h	<Twtr> clocks	Fh-Ch	Reserved		
Bits	Description										
3h-0h	Reserved										
Bh-4h	<Twtr> clocks										
Fh-Ch	Reserved										
7:5	Reserved.										
4:0	Tcwl: CAS write latency. Read-write. Reset: 5h. Specifies the number of memory clock cycles from internal write command to first write data in at the DRAM. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>04h-00h</td> <td>Reserved</td> </tr> <tr> <td>11h-05h</td> <td><Tcwl> clocks</td> </tr> <tr> <td>1Fh-12h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	04h-00h	Reserved	11h-05h	<Tcwl> clocks	1Fh-12h	Reserved		
Bits	Description										
04h-00h	Reserved										
11h-05h	<Tcwl> clocks										
1Fh-12h	Reserved										

D18F2x210_dct[1:0]_nbp[3:0] DRAM NB P-state

See 2.9.1 [DCT Configuration Registers]. For D18F2x210_dct[1:0]_nbp[x], x=D18F1x10C[NbPsSel]; see D18F1x10C[NbPsSel].

Table 177: BIOS Recommendations for RdPtrInit

Condition		D18F2x210_dct[1:0]_nbp[3:0]
NCLK:MCLK ratio	DDR rate (MT/s)	RdPtrInit
< 2:1	< 2133	0011b
< 2:1	2133 <= rate <= 2400	0011b or 0010b See Note 1.
>=2:1	< 1866	0110b
>=2:1	1866 <= rate < 2400	0101b
>=2:1	2400	0100b

1. For each NB P-state, IF any D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[MemPstate][WrDat-GrossDly] ==0 THEN RdPtrInit=0010b ELSE RdPtrInit=0011b.

Bits	Description																		
31:22	<p>MaxRdLatency: maximum read latency. Read-write. Reset: 000h. BIOS: 2.9.5.9.5. Specifies the maximum round-trip latency in the system from the processor to the DRAM devices and back. The DRAM controller uses this to help determine when the first two beats of incoming DRAM read data can be safely transferred to the NCLK domain. The time includes the asynchronous and synchronous latencies.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000h</td> <td>0 NCLKs</td> </tr> <tr> <td>3FEh-001h</td> <td><MaxRdLatency> NCLKs</td> </tr> <tr> <td>3FFh</td> <td>1023 NCLKs</td> </tr> </tbody> </table>	Bits	Description	000h	0 NCLKs	3FEh-001h	<MaxRdLatency> NCLKs	3FFh	1023 NCLKs										
Bits	Description																		
000h	0 NCLKs																		
3FEh-001h	<MaxRdLatency> NCLKs																		
3FFh	1023 NCLKs																		
21:19	Reserved.																		
18:16	<p>DataTxFifoWrDly: data transmit FIFO write delay. Read-write. Reset: 0. BIOS: 0h. Specifies the DCT to phy write data FIFO delay. BIOS must program this field for the current or target NB P-state prior to a frequency change or NB P-state change. See also 2.9.5.2 [NB P-state Specific Configuration].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0 MEMCLK</td> </tr> <tr> <td>001b</td> <td>0.5 MEMCLK</td> </tr> <tr> <td>010b</td> <td>1.0 MEMCLK</td> </tr> <tr> <td>011b</td> <td>1.5 MEMCLKs</td> </tr> <tr> <td>100b</td> <td>2.0 MEMCLKs</td> </tr> <tr> <td>101b</td> <td>2.5 MEMCLKs</td> </tr> <tr> <td>110b</td> <td>3.0 MEMCLKs</td> </tr> <tr> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	0 MEMCLK	001b	0.5 MEMCLK	010b	1.0 MEMCLK	011b	1.5 MEMCLKs	100b	2.0 MEMCLKs	101b	2.5 MEMCLKs	110b	3.0 MEMCLKs	111b	Reserved
Bits	Description																		
000b	0 MEMCLK																		
001b	0.5 MEMCLK																		
010b	1.0 MEMCLK																		
011b	1.5 MEMCLKs																		
100b	2.0 MEMCLKs																		
101b	2.5 MEMCLKs																		
110b	3.0 MEMCLKs																		
111b	Reserved																		
15:4	Reserved.																		
3:0	<p>RdPtrInit: read pointer initial value. Read-write. Reset: 6h. BIOS: Table 177. There is a synchronization FIFO between the NB clock domain and memory clock domain. Each increment of this field positions the read pointer one half clock cycle closer to the write pointer thereby reducing the latency through the FIFO. BIOS must program this field for the current or target NB P-state prior to a frequency change or NB P-state change. See also 2.9.5.2 [NB P-state Specific Configuration].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0001b-0000b</td> <td>Reserved</td> </tr> <tr> <td>0010b</td> <td>3 MEMCLKs</td> </tr> <tr> <td>0011b</td> <td>2.5 MEMCLKs</td> </tr> <tr> <td>0100b</td> <td>2 MEMCLKs</td> </tr> <tr> <td>0101b</td> <td>1.5 MEMCLKs</td> </tr> <tr> <td>0110b</td> <td>1 MEMCLK</td> </tr> <tr> <td>1111b-0111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0001b-0000b	Reserved	0010b	3 MEMCLKs	0011b	2.5 MEMCLKs	0100b	2 MEMCLKs	0101b	1.5 MEMCLKs	0110b	1 MEMCLK	1111b-0111b	Reserved		
Bits	Description																		
0001b-0000b	Reserved																		
0010b	3 MEMCLKs																		
0011b	2.5 MEMCLKs																		
0100b	2 MEMCLKs																		
0101b	1.5 MEMCLKs																		
0110b	1 MEMCLK																		
1111b-0111b	Reserved																		

D18F2x214_dct[1:0]_mp[1:0] DRAM Timing 4

Reset: 0001_0202h.

Bits	Description
31:20	Reserved.

19:16	<p>TwrrwrSdSc: write to write timing same DIMM same chip select. Read-write. BIOS: See 2.9.5.6.2 [TwrrwrSdSc, TwrrwrSdDc, TwrrwrDd (Write to Write Timing)]. Specifies the minimum number of cycles from the last clock of virtual CAS of the first write-burst operation to the clock in which CAS is asserted for a following write-burst operation.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved</td> </tr> <tr> <td>1h</td> <td>1 clock</td> </tr> <tr> <td>Ah-2h</td> <td><TwrrwrSdSc> clocks</td> </tr> <tr> <td>Bh</td> <td>11 clocks</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	Reserved	1h	1 clock	Ah-2h	<TwrrwrSdSc> clocks	Bh	11 clocks	Fh-Ch	Reserved
<u>Bits</u>	<u>Description</u>												
0h	Reserved												
1h	1 clock												
Ah-2h	<TwrrwrSdSc> clocks												
Bh	11 clocks												
Fh-Ch	Reserved												
15:12	Reserved.												
11:8	TwrrwrSdDc: write to write timing same DIMM different chip select. See: TwrrwrDd.												
7:4	Reserved.												
3:0	<p>TwrrwrDd: write to write timing different DIMM. Read-write. BIOS: See 2.9.5.6.2 [TwrrwrSdSc, TwrrwrSdDc, TwrrwrDd (Write to Write Timing)]. Specifies the minimum number of cycles from the last clock of virtual CAS of the first write-burst operation to the clock in which CAS is asserted for a following write-burst operation.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>1h-0h</td> <td>Reserved</td> </tr> <tr> <td>Bh-2h</td> <td><TwrrwrDd> clocks</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	1h-0h	Reserved	Bh-2h	<TwrrwrDd> clocks	Fh-Ch	Reserved				
<u>Bits</u>	<u>Description</u>												
1h-0h	Reserved												
Bh-2h	<TwrrwrDd> clocks												
Fh-Ch	Reserved												

D18F2x218_dct[1:0]_mp[1:0] DRAM Timing 5

Reset: 0103_0203h. See [2.9.1 \[DCT Configuration Registers\]](#).

Bits	Description								
31:28	Reserved.								
27:24	<p>TrdrdSdSc: read to read timing same DIMM same chip select. Read-write. BIOS: See 2.9.5.6.1 [TrdrdSdSc, TrdrdSdDd, and TrdrdDd (Read to Read Timing)]. Specifies the minimum number of cycles from the last clock of virtual CAS of a first read-burst operation to the clock in which CAS is asserted for a following read-burst operation.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved</td> </tr> <tr> <td>Bh-1h</td> <td><TrdrdSdSc> clocks</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	Reserved	Bh-1h	<TrdrdSdSc> clocks	Fh-Ch	Reserved
<u>Bits</u>	<u>Description</u>								
0h	Reserved								
Bh-1h	<TrdrdSdSc> clocks								
Fh-Ch	Reserved								
23:20	Reserved.								
19:16	TrdrdSdDc: read to read timing same DIMM different chip select. See: TrdrdDd.								
15:12	Reserved.								

11:8	<p>Twrrd: write to read DIMM termination turnaround. Read-write. BIOS: See 2.9.5.6.3 [Twrrd (Write to Read DIMM Termination Turn-around)]. Specifies the minimum number of cycles from the last clock of virtual CAS of the first write-burst operation to the clock in which CAS is asserted for a following read-burst operation, both to different chip selects.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved</td> </tr> <tr> <td>Bh-1h</td> <td><Twrrd> clocks</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	Reserved	Bh-1h	<Twrrd> clocks	Fh-Ch	Reserved
<u>Bits</u>	<u>Description</u>								
0h	Reserved								
Bh-1h	<Twrrd> clocks								
Fh-Ch	Reserved								
7:4	Reserved.								
3:0	<p>TrdrdDd: read to read timing different DIMM. Read-write. BIOS: See 2.9.5.6.1 [TrdrdSdSc, TrdrdSdDd, and TrdrdDd (Read to Read Timing)]. Specifies the minimum number of cycles from the last clock of virtual CAS of a first read-burst operation to the clock in which CAS is asserted for a following read-burst operation.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>1h-0h</td> <td>Reserved</td> </tr> <tr> <td>Bh-2h</td> <td><TrdrdDd> clocks</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	1h-0h	Reserved	Bh-2h	<TrdrdDd> clocks	Fh-Ch	Reserved
<u>Bits</u>	<u>Description</u>								
1h-0h	Reserved								
Bh-2h	<TrdrdDd> clocks								
Fh-Ch	Reserved								

D18F2x21C_dct[1:0]_mp[1:0] DRAM Timing 6

Reset: 0004_0300h. See [2.9.1 \[DCT Configuration Registers\]](#).

Bits	Description								
31:21	Reserved.								
20:16	<p>TrwtWB: read to write turnaround for opportunistic write bursting. Read-write. BIOS: TrwtTO + 1. Specifies the minimum number of clock cycles from the last clock of virtual CAS of a first read-burst operation to the clock in which CAS is asserted for a following write-burst operation.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>02h-00h</td> <td>Reserved</td> </tr> <tr> <td>17h-03h</td> <td><TrwtWB> clocks</td> </tr> <tr> <td>1Fh-18h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	02h-00h	Reserved	17h-03h	<TrwtWB> clocks	1Fh-18h	Reserved
<u>Bits</u>	<u>Description</u>								
02h-00h	Reserved								
17h-03h	<TrwtWB> clocks								
1Fh-18h	Reserved								
15:13	Reserved.								
12:8	<p>TrwtTO: read to write turnaround. Read-write. BIOS: See 2.9.5.6.4 [TrwtTO (Read-to-Write Turnaround for Data, DQS Contention)]. Specifies the minimum number of clock cycles from the last clock of virtual CAS of a first read-burst operation to the clock in which CAS is asserted for a following write-burst operation.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>01h-00h</td> <td>Reserved</td> </tr> <tr> <td>16h-02h</td> <td><TrwtTO> clocks</td> </tr> <tr> <td>1Fh-17h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	01h-00h	Reserved	16h-02h	<TrwtTO> clocks	1Fh-17h	Reserved
<u>Bits</u>	<u>Description</u>								
01h-00h	Reserved								
16h-02h	<TrwtTO> clocks								
1Fh-17h	Reserved								
7:0	Reserved.								

D18F2x220_dct[1:0] DRAM Timing 7

Reset: 0000_0C04h. See [2.9.1 \[DCT Configuration Registers\]](#).

Bits	Description								
31:13	Reserved.								
12:8	<p>Tmod: mode register command delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time in memory clock cycles from an MRS command to another non-MRS command (excluding NOP and DES), all to the same chip select.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1h-0h</td> <td>Reserved</td> </tr> <tr> <td>14h-2h</td> <td><Tmod> clocks</td> </tr> <tr> <td>1Fh-15h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	1h-0h	Reserved	14h-2h	<Tmod> clocks	1Fh-15h	Reserved
Bits	Description								
1h-0h	Reserved								
14h-2h	<Tmod> clocks								
1Fh-15h	Reserved								
7:4	Reserved.								
3:0	<p>Tmrd: mode register command cycle time. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time in memory clock cycles from an MRS command to another MRS command, all to the same chip select.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1h-0h</td> <td>Reserved</td> </tr> <tr> <td>8h-2h</td> <td><Tmrd> clocks</td> </tr> <tr> <td>Fh-9h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	1h-0h	Reserved	8h-2h	<Tmrd> clocks	Fh-9h	Reserved
Bits	Description								
1h-0h	Reserved								
8h-2h	<Tmrd> clocks								
Fh-9h	Reserved								

D18F2x224_dct[1:0] DRAM Timing 8

Reset: 0000_0408h. See 2.9.1 [DCT Configuration Registers].

Bits	Description																																				
31:11	Reserved.																																				
10:8	<p>Tzqcs: Zq short cal command delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time in memory clock cycles from a ZQCS command to any other command (excluding NOP and DES) on the channel.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Reserved</td> <td>100b</td> <td>64 clocks</td> </tr> <tr> <td>001b</td> <td>16 clocks</td> <td>101b</td> <td>80 clocks</td> </tr> <tr> <td>010b</td> <td>32 clocks</td> <td>110b</td> <td>96 clocks</td> </tr> <tr> <td>011b</td> <td>48 clocks</td> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	Reserved	100b	64 clocks	001b	16 clocks	101b	80 clocks	010b	32 clocks	110b	96 clocks	011b	48 clocks	111b	Reserved																
Bits	Description	Bits	Description																																		
000b	Reserved	100b	64 clocks																																		
001b	16 clocks	101b	80 clocks																																		
010b	32 clocks	110b	96 clocks																																		
011b	48 clocks	111b	Reserved																																		
7:4	Reserved.																																				
3:0	<p>Tzqoper: Zq long cal command delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time in memory clock cycles from a ZQCL command to any other command (excluding NOP and DES) on the channel.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>Reserved</td> <td>1000b</td> <td>256 clocks</td> </tr> <tr> <td>0001b</td> <td>32 clocks</td> <td>1001b</td> <td>288 clocks</td> </tr> <tr> <td>0010b</td> <td>64 clocks</td> <td>1010b</td> <td>320 clocks</td> </tr> <tr> <td>0011b</td> <td>96 clocks</td> <td>1011b</td> <td>352 clocks</td> </tr> <tr> <td>0100b</td> <td>128 clocks</td> <td>1100b</td> <td>384 clocks</td> </tr> <tr> <td>0101b</td> <td>160 clocks</td> <td>1111b-1101b</td> <td>Reserved</td> </tr> <tr> <td>0110b</td> <td>192 clocks</td> <td></td> <td></td> </tr> <tr> <td>0111b</td> <td>224 clocks</td> <td></td> <td></td> </tr> </tbody> </table>	Bits	Description	Bits	Description	0000b	Reserved	1000b	256 clocks	0001b	32 clocks	1001b	288 clocks	0010b	64 clocks	1010b	320 clocks	0011b	96 clocks	1011b	352 clocks	0100b	128 clocks	1100b	384 clocks	0101b	160 clocks	1111b-1101b	Reserved	0110b	192 clocks			0111b	224 clocks		
Bits	Description	Bits	Description																																		
0000b	Reserved	1000b	256 clocks																																		
0001b	32 clocks	1001b	288 clocks																																		
0010b	64 clocks	1010b	320 clocks																																		
0011b	96 clocks	1011b	352 clocks																																		
0100b	128 clocks	1100b	384 clocks																																		
0101b	160 clocks	1111b-1101b	Reserved																																		
0110b	192 clocks																																				
0111b	224 clocks																																				

D18F2x228_dct[1:0] DRAM Timing 9

See 2.9.1 [DCT Configuration Registers].

Bits	Description						
31:24	Tstag3: auto refresh stagger time for logical DIMM 3. See: Tstag0.						
23:16	Tstag2: auto refresh stagger time for logical DIMM 2. See: Tstag0.						
15:8	Tstag1: auto refresh stagger time for logical DIMM 1. See: Tstag0.						
7:0	<p>Tstag0: auto refresh stagger time for logical DIMM 0. Read-write. Reset: 00h. BIOS: 14h. Specifies the number of clocks between auto refresh commands to different ranks of a DIMM when D18F2x90_dct[1:0][StagRefEn]=1 or D18F2x[6C:60]_dct[1:0][RankDef] != 0. BitsDescription</p> <table> <tr> <td>00h</td> <td>0 clocks</td> </tr> <tr> <td>FEh-01h</td> <td><Tstag0> clocks</td> </tr> <tr> <td>FFh</td> <td>255 clocks</td> </tr> </table>	00h	0 clocks	FEh-01h	<Tstag0> clocks	FFh	255 clocks
00h	0 clocks						
FEh-01h	<Tstag0> clocks						
FFh	255 clocks						

D18F2x22C_dct[1:0]_mp[1:0] DRAM Timing 10

Reset: 0000_000Ch. See 2.9.1 [DCT Configuration Registers].

Bits	Description																												
31:5	Reserved.																												
4:0	<p>Twr: write recovery. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time from the last data write until the chip select bank precharge.</p> <table> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>4h-0h</td> <td>Reserved</td> </tr> <tr> <td>8h-5h</td> <td>8 to 5 clocks</td> </tr> <tr> <td>9h</td> <td>Reserved</td> </tr> <tr> <td>Ah</td> <td>10 clocks</td> </tr> <tr> <td>Bh</td> <td>Reserved</td> </tr> <tr> <td>Ch</td> <td>12 clocks</td> </tr> <tr> <td>Dh</td> <td>Reserved</td> </tr> <tr> <td>Eh</td> <td>14 clocks</td> </tr> <tr> <td>Fh</td> <td>Reserved</td> </tr> <tr> <td>10h</td> <td>16 clocks</td> </tr> <tr> <td>11h</td> <td>Reserved</td> </tr> <tr> <td>12h</td> <td>18 clocks</td> </tr> <tr> <td>1Fh-13h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	4h-0h	Reserved	8h-5h	8 to 5 clocks	9h	Reserved	Ah	10 clocks	Bh	Reserved	Ch	12 clocks	Dh	Reserved	Eh	14 clocks	Fh	Reserved	10h	16 clocks	11h	Reserved	12h	18 clocks	1Fh-13h	Reserved
<u>Bits</u>	<u>Description</u>																												
4h-0h	Reserved																												
8h-5h	8 to 5 clocks																												
9h	Reserved																												
Ah	10 clocks																												
Bh	Reserved																												
Ch	12 clocks																												
Dh	Reserved																												
Eh	14 clocks																												
Fh	Reserved																												
10h	16 clocks																												
11h	Reserved																												
12h	18 clocks																												
1Fh-13h	Reserved																												

D18F2x[234:230]_dct[1:0] DRAM Read ODT Pattern [High:Low]

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. This register is used by BIOS to specify the state of the ODT pins during DDR reads. F2x230 is used to control chip selects 0-3. F2x234 is used to control chip selects 4-7. See section 2.9.5.6.5 [DRAM ODT Control] for more information.

Bits	Description
31:28	Reserved.

27:24	RdOdtPatCs73: read ODT pattern chip select [7,3] . See: RdOdtPatCs40.
23:20	Reserved.
19:16	RdOdtPatCs62: read ODT pattern chip select [6,2] . See: RdOdtPatCs40.
15:12	Reserved.
11:8	RdOdtPatCs51: read ODT pattern chip select [5,1] . See: RdOdtPatCs40.
7:4	Reserved.
3:0	RdOdtPatCs40: read ODT pattern chip select [4,0] . Read-write. Specifies the state of ODT[3:0] pins when a read occurs to the specified chip select.

D18F2x[23C:238]_dct[1:0] DRAM Write ODT Pattern [High:Low]

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. This register is used by BIOS to specify the state of the ODT pins during DDR writes. F2x238 is used to control chip selects 0-3. F2x23C is used to control chip selects 4-7. See section 2.9.5.6.5 [DRAM ODT Control] for more information.

Bits	Description
31:28	Reserved.
27:24	WrOdtPatCs73: write ODT pattern chip select [7,3] . See: WrOdtPatCs40.
23:20	Reserved.
19:16	WrOdtPatCs62: write ODT pattern chip select [6,2] . See: WrOdtPatCs40.
15:12	Reserved.
11:8	WrOdtPatCs51: write ODT pattern chip select [5,1] . See: WrOdtPatCs40.
7:4	Reserved.
3:0	WrOdtPatCs40: write ODT pattern chip select [4,0] . Read-write. Specifies the state of ODT[3:0] pins when a write occurs to the specified chip select.

D18F2x240_dct[1:0]_mp[1:0] DRAM ODT Control

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers].

Bits	Description						
31:15	Reserved.						
14:12	WrOdtOnDuration: write ODT on duration. Read-write. BIOS: 6. Specifies the number of memory clock cycles that ODT is asserted for an eight-beat write burst. The controller will shorten the ODT pulse duration by two clock cycles if the burst is chopped. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>5h-0h</td> <td>Reserved</td> </tr> <tr> <td>7h-6h</td> <td><WrOdtOnDuration> clocks</td> </tr> </tbody> </table>	Bits	Description	5h-0h	Reserved	7h-6h	<WrOdtOnDuration> clocks
Bits	Description						
5h-0h	Reserved						
7h-6h	<WrOdtOnDuration> clocks						
11	Reserved.						
10:8	WrOdtTrnOnDly: Write ODT Turn On Delay. Read-write. BIOS: 0. Specifies the number of memory clock cycles that ODT assertion is delayed relative to write CAS. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0 clocks</td> </tr> <tr> <td>7h-1h</td> <td><WrOdtTrnOnDly> clocks, Reserved if (WrOdtOnDuration=0)</td> </tr> </tbody> </table>	Bits	Description	0h	0 clocks	7h-1h	<WrOdtTrnOnDly> clocks, Reserved if (WrOdtOnDuration=0)
Bits	Description						
0h	0 clocks						
7h-1h	<WrOdtTrnOnDly> clocks, Reserved if (WrOdtOnDuration=0)						
7	Reserved.						

6:4	RdOdtOnDuration: Read ODT On Duration. Read-write. BIOS: 6. Specifies the number of memory clock cycles that ODT is asserted for an eight-beat read burst. The controller will shorten the ODT pulse duration by two clock cycles if the burst is chopped. <u>Bits</u> <u>Description</u> 5h-0h Reserved 7h-6h <RdOdtOnDuration> clocks
3:0	RdOdtTrnOnDly: Read ODT Turn On Delay. Read-write. BIOS: MAX(0, D18F2x200_dct[1:0]_mp[1:0][Tcl] - D18F2x20C_dct[1:0]_mp[1:0][Tcwl]). Specifies the number of clock cycles that ODT assertion is delayed relative to read CAS. <u>Bits</u> <u>Description</u> 0h 0 clocks Fh-0h <RdOdtTrnOnDly> clocks, Reserved if (RdOdtOnDuration=0)

D18F2x244_dct[1:0] DRAM Controller Miscellaneous 3

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:4	Reserved.
3:0	PrtlChPDDynDly: partial channel power down dynamic delay. Read-write. BIOS: IF(D18F2xA8_dct[1:0][PrtlChPDEnhEn]) THEN 4h ELSE 0 ENDIF. Specifies the channel idle hysteresis for fast exit/slow exit mode changes when D18F2xA8_dct[1:0][PrtlChPDEnhEn]=1. <u>Bits</u> <u>Description</u> 0h 0 clocks 7h-1h <PrtlChPDDynDly*32> clocks 8h 256 clocks Fh-9h Reserved

D18F2x248_dct[1:0]_mp[1:0] DRAM Power Management 0

Reset: 0000_0A03h. See 2.9.1 [DCT Configuration Registers].

Bits	Description
31	RxChMntClkEn: Receive channel maintenance clocks. Read-Write. BIOS: See 2.9.5.11. 1=Enable receive channel maintenance clocks to improve internal timing margin at the cost of some extra power. 0=Disable clocks. To disable clocks, BIOS must first disable clock generation in the phy (see D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0][RxSsbMntClkEn]).
30	Reserved.
29:24	AggrPDDelay: aggressive power down delay. Read-Write. BIOS: 20h. Specifies a hysteresis count from the last DRAM activity for the DCT to close pages prior to precharge power down. Reserved if ~D18F2xA8_dct[1:0][AggrPDEn]. See PchgPDEnDelay and D18F2x94_dct[1:0][PowerDownEn]. <u>Bits</u> <u>Description</u> 00h 64 clocks 01h 1 clock 3Eh-02h <AggrPDDelay> clocks 3Fh 63 clocks
23:22	Reserved.

21:16	<p>PchgPDEnDelay: precharge power down entry delay. Read-write. BIOS: $\text{MAX}(\text{D18F2x200_dct}[1:0]_mp[1:0][\text{Tcl}] + 5, \text{D18F2x20C_dct}[1:0]_mp[1:0][\text{Tcw}] + \text{D18F2x22C_dct}[1:0]_mp[1:0][\text{Twr}] + 5, \text{D18F2x220_dct}[1:0][\text{Tmod}])$. Specifies the power down entry delay. If $\text{D18F2xA8_dct}[1:0][\text{AggrPDEn}] = 0$, this delay behaves as a hysteresis. This field must satisfy the minimum power down entry delay requirements. See also D18F2x94_dct[1:0][PowerDownEn].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>64 clocks</td> </tr> <tr> <td>01h</td> <td>1 clock</td> </tr> <tr> <td>3Eh-02h</td> <td><PchgPDEnDelay> clocks</td> </tr> <tr> <td>3Fh</td> <td>63 clocks</td> </tr> </tbody> </table>	Bits	Description	00h	64 clocks	01h	1 clock	3Eh-02h	<PchgPDEnDelay> clocks	3Fh	63 clocks
Bits	Description										
00h	64 clocks										
01h	1 clock										
3Eh-02h	<PchgPDEnDelay> clocks										
3Fh	63 clocks										
15:13	Reserved.										
12:8	<p>Txpdll: exit DLL and precharge powerdown to command delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time that the DCT waits to issue a command after exiting precharge power-down mode if the DLL was also disabled.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>09h-00h</td> <td>Reserved</td> </tr> <tr> <td>1Dh-0Ah</td> <td><Txpdll> clocks</td> </tr> <tr> <td>1Fh-1Eh</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	09h-00h	Reserved	1Dh-0Ah	<Txpdll> clocks	1Fh-1Eh	Reserved		
Bits	Description										
09h-00h	Reserved										
1Dh-0Ah	<Txpdll> clocks										
1Fh-1Eh	Reserved										
7:4	Reserved.										
3:0	<p>Txp: exit precharge PD to command delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time that the DCT waits to issue a command after exiting precharge powerdown mode.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>2h-0h</td> <td>Reserved</td> </tr> <tr> <td>8h-3h</td> <td><Txp> clocks</td> </tr> <tr> <td>Fh-9h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	2h-0h	Reserved	8h-3h	<Txp> clocks	Fh-9h	Reserved		
Bits	Description										
2h-0h	Reserved										
8h-3h	<Txp> clocks										
Fh-9h	Reserved										

D18F2x24C_dct[1:0] DRAM Power Management 1

Reset: 0214_0803h. See [2.9.1 \[DCT Configuration Registers\]](#).

Bits	Description								
31:30	Reserved.								
29:24	<p>Tksrx: clock stable to self refresh exit delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time in memory clock cycles that the DCT waits to assert CKE after clock frequency is stable.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>01h-00h</td> <td>Reserved</td> </tr> <tr> <td>0Eh-02h</td> <td><Tksrx> clocks</td> </tr> <tr> <td>3Fh-0Fh</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	01h-00h	Reserved	0Eh-02h	<Tksrx> clocks	3Fh-0Fh	Reserved
Bits	Description								
01h-00h	Reserved								
0Eh-02h	<Tksrx> clocks								
3Fh-0Fh	Reserved								
23:22	Reserved.								
21:16	<p>Tksre: self refresh to command delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time in memory clock cycles that the DCT waits to remove external clocks after entering self refresh or powerdown.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>04h-00h</td> <td>Reserved</td> </tr> <tr> <td>27h-05h</td> <td><Tksre> clocks</td> </tr> <tr> <td>3Fh-28h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	04h-00h	Reserved	27h-05h	<Tksre> clocks	3Fh-28h	Reserved
Bits	Description								
04h-00h	Reserved								
27h-05h	<Tksre> clocks								
3Fh-28h	Reserved								
15:14	Reserved.								

13:8	Tckesr: self refresh to command delay. Read-write. BIOS: See 2.9.5.5. Specifies the minimum time in memory clock cycles that the DCT waits to issue a command after entering self refresh. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>01h-00h</td> <td>Reserved</td> </tr> <tr> <td>2Bh-02h</td> <td><Tckesr> clocks</td> </tr> <tr> <td>3Fh-2Ch</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	01h-00h	Reserved	2Bh-02h	<Tckesr> clocks	3Fh-2Ch	Reserved
Bits	Description								
01h-00h	Reserved								
2Bh-02h	<Tckesr> clocks								
3Fh-2Ch	Reserved								
7:4	Reserved.								
3:0	Tpd: minimum power down entry to exit. Read-write. BIOS: See 2.9.5.5. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved</td> </tr> <tr> <td>Ah-1h</td> <td><Tpd> clocks</td> </tr> <tr> <td>Fh-Bh</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	Reserved	Ah-1h	<Tpd> clocks	Fh-Bh	Reserved
Bits	Description								
0h	Reserved								
Ah-1h	<Tpd> clocks								
Fh-Bh	Reserved								

D18F2x250_dct[1:0] DRAM Loopback and Training Control

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description										
31:13	Reserved.										
12	CmdSendInProg: command in progress. Read-only. 0=DCT is idle. 1=DCT is busy.										
11	SendCmd: send command. Read-write. 0=Stop command generation. 1=Begin command generation as specified in CmdTgt, CmdType, and D18F2x260_dct[1:0][CmdCount]. BIOS must set this field to a 0 after a command series is completed. Reserved if ~CmdTestEnable.										
10	TestStatus: test status. Read-only. 0=Command generation is in progress. 1=Command generation has completed. Reserved if ~(SendCmd & (D18F2x260_dct[1:0][CmdCount] > 0 StopOnErr)).										
9:8	CmdTgt: command target. Read-write. Specifies the SendCmd command target address mode. See D18F2x25[8,4]_dct[1:0]. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Issue commands to address Target A</td> </tr> <tr> <td>01b</td> <td>Issue alternating commands to address Target A and Target B</td> </tr> <tr> <td>11b-10b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	Issue commands to address Target A	01b	Issue alternating commands to address Target A and Target B	11b-10b	Reserved		
Bits	Description										
00b	Issue commands to address Target A										
01b	Issue alternating commands to address Target A and Target B										
11b-10b	Reserved										
7:5	CmdType: command type. Read-write. Specifies the SendCmd command type. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Read</td> </tr> <tr> <td>001b</td> <td>Write</td> </tr> <tr> <td>010b</td> <td>Alternating write and read</td> </tr> <tr> <td>111b-011b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	Read	001b	Write	010b	Alternating write and read	111b-011b	Reserved
Bits	Description										
000b	Read										
001b	Write										
010b	Alternating write and read										
111b-011b	Reserved										
4	StopOnErr: stop on error. Read-write. Specifies the DCT behavior if a data comparison error occurs. 1=Stop command generation. 0=Continue command generation. If StopOnErr=1, BIOS must program ResetAllErr=1 when programming SendCmd=1.										
3	ResetAllErr: reset all errors. Read; write-1-only; cleared-by-hardware. 1=Clear error status bits and error counters in D18F2x264_dct[1:0], D18F2x268_dct[1:0], and D18F2x268_dct[1:0].										
2	CmdTestEnable: command test enable. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. 0=Disable the command generation mode. 1=Enable the command generation mode. See SendCmd.										
1:0	Reserved.										

D18F2x25[8,4]_dct[1:0] DRAM Target [B, A] Base

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description
31:27	Reserved.
26:24	TgtChipSelect: target chip select. Read-write. Specifies the chip select. 0=CS0, 1=CS1, ..., 7=CS7.
23:21	TgtBank: target bank [2:0]. Read-write. Specifies the bank address.
20:10	Reserved.
9:0	TgtAddress[9:0]: target address [9:0]. Read-write. Specifies the column address bits [9:0]. The corresponding address sequence in a command series is as follows: TgtAddress[9:3] is incremented by one, with wrap around, after each command if <code>D18F2x250_dct[1:0][CmdType] = 00xb</code> or if <code>D18F2x250_dct[1:0][CmdType] = 010b</code> and <code>D18F2x250_dct[1:0][CmdTgt] = 01b</code> . TgtAddress[9:3] is incremented by one, with wrap around, after each command pair if <code>D18F2x250_dct[1:0][CmdType] = 010b</code> and <code>D18F2x250_dct[1:0][CmdTgt] = 00b</code> .

D18F2x260_dct[1:0] DRAM Command 1

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description						
31:21	Reserved.						
20:0	CmdCount: command count. Read-write. Specifies the maximum number of commands to generate when <code>D18F2x250_dct[1:0][SendCmd]=1</code> . See also <code>D18F2x250_dct[1:0][StopOnErr]</code> . <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Infinite commands</td> </tr> <tr> <td>1F_FFFFh-1h</td> <td><CmdCount> commands</td> </tr> </tbody> </table>	Bits	Description	0h	Infinite commands	1F_FFFFh-1h	<CmdCount> commands
Bits	Description						
0h	Infinite commands						
1F_FFFFh-1h	<CmdCount> commands						

D18F2x264_dct[1:0] DRAM Status 0

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description
------	-------------

31:25	<p>ErrDqNum: error DQ number. Read-only. Indicates the DQ bit of the first error occurrence when <code>D18F2x264_dct[1:0][ErrCnt] > 0</code>. Cleared by <code>D18F2x250_dct[1:0][ResetAllErr]</code>.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Data[0]</td> </tr> <tr> <td>3Eh-01h</td> <td>Data[<ErrDqNum>]</td> </tr> <tr> <td>3Fh</td> <td>Data[63]</td> </tr> <tr> <td>40h</td> <td>ECC[0]</td> </tr> <tr> <td>46h-41h</td> <td>ECC[<ErrDqNum>-40h]</td> </tr> <tr> <td>47h</td> <td>ECC[7]</td> </tr> <tr> <td>7Fh-48h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00h	Data[0]	3Eh-01h	Data[<ErrDqNum>]	3Fh	Data[63]	40h	ECC[0]	46h-41h	ECC[<ErrDqNum>-40h]	47h	ECC[7]	7Fh-48h	Reserved
Bits	Description																
00h	Data[0]																
3Eh-01h	Data[<ErrDqNum>]																
3Fh	Data[63]																
40h	ECC[0]																
46h-41h	ECC[<ErrDqNum>-40h]																
47h	ECC[7]																
7Fh-48h	Reserved																
24:0	<p>ErrCnt: error count. Read; set-by-hardware; write-1-to-clear. Specifies a saturating counter indicating the number of DQ bit errors detected. Counts a maximum of 72 errors per bit-time. Status is accumulated until cleared by <code>D18F2x250_dct[1:0][ResetAllErr]</code>. Errors can an be masked on per-bit basis by programming <code>D18F2x274_dct[1:0]</code> and <code>D18F2x278_dct[1:0]</code>.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0 errors</td> </tr> <tr> <td>1FF_FFFDh-1h</td> <td><ErrCnt> errors</td> </tr> <tr> <td>1FF_FFFEh</td> <td>1FF_FFFEh errors</td> </tr> <tr> <td>1FF_FFFFh</td> <td>1FF_FFFFh or more errors</td> </tr> </tbody> </table>	Bits	Description	0h	0 errors	1FF_FFFDh-1h	<ErrCnt> errors	1FF_FFFEh	1FF_FFFEh errors	1FF_FFFFh	1FF_FFFFh or more errors						
Bits	Description																
0h	0 errors																
1FF_FFFDh-1h	<ErrCnt> errors																
1FF_FFFEh	1FF_FFFEh errors																
1FF_FFFFh	1FF_FFFFh or more errors																

D18F2x268_dct[1:0] DRAM Status 1

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description														
31:18	Reserved.														
17:0	<p>NibbleErrSts: nibble error status. Read-only. Indicates error detection status on a per nibble basis when <code>D18F2x264_dct[1:0][ErrCnt] > 0</code>. Status is accumulated until cleared by <code>D18F2x250_dct[1:0][ResetAllErr]</code>.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Data[3:0]</td> </tr> <tr> <td>[1]</td> <td>Data[7:4]</td> </tr> <tr> <td>[14:2]</td> <td>Data[(<NibbleErrSts>*4)+3:<NibbleErrSts>*4]</td> </tr> <tr> <td>[15]</td> <td>Data[63:60]</td> </tr> <tr> <td>[16]</td> <td>ECC[3:0]</td> </tr> <tr> <td>[17]</td> <td>ECC[7:4]</td> </tr> </tbody> </table>	Bit	Description	[0]	Data[3:0]	[1]	Data[7:4]	[14:2]	Data[(<NibbleErrSts>*4)+3:<NibbleErrSts>*4]	[15]	Data[63:60]	[16]	ECC[3:0]	[17]	ECC[7:4]
Bit	Description														
[0]	Data[3:0]														
[1]	Data[7:4]														
[14:2]	Data[(<NibbleErrSts>*4)+3:<NibbleErrSts>*4]														
[15]	Data[63:60]														
[16]	ECC[3:0]														
[17]	ECC[7:4]														

D18F2x26C_dct[1:0] DRAM Status 2

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description
------	-------------

31:18	Reserved.														
17:0	<p>NibbleErr180Sts: nibble error 180 status. Read-only. Indicates error detection status on a per nibble basis when <code>D18F2x264_dct[1:0][ErrCnt] > 0</code>, comparing read data against data shifted 1-bit time earlier. Status is accumulated until cleared by <code>D18F2x250_dct[1:0][ResetAllErr]</code>.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Data[3:0]</td> </tr> <tr> <td>[1]</td> <td>Data[7:4]</td> </tr> <tr> <td>[14:2]</td> <td>Data[(<code><NibbleErr180Sts>*4</code>)+3:<code><NibbleErr180Sts>*4</code>]</td> </tr> <tr> <td>[15]</td> <td>Data[63:60]</td> </tr> <tr> <td>[16]</td> <td>ECC[3:0]</td> </tr> <tr> <td>[17]</td> <td>ECC[7:4]</td> </tr> </tbody> </table>	Bit	Description	[0]	Data[3:0]	[1]	Data[7:4]	[14:2]	Data[(<code><NibbleErr180Sts>*4</code>)+3: <code><NibbleErr180Sts>*4</code>]	[15]	Data[63:60]	[16]	ECC[3:0]	[17]	ECC[7:4]
Bit	Description														
[0]	Data[3:0]														
[1]	Data[7:4]														
[14:2]	Data[(<code><NibbleErr180Sts>*4</code>)+3: <code><NibbleErr180Sts>*4</code>]														
[15]	Data[63:60]														
[16]	ECC[3:0]														
[17]	ECC[7:4]														

D18F2x270_dct[1:0] DRAM PRBS

See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description										
31	Reserved.										
30:24	Reserved.										
23:19	Reserved.										
18:0	<p>DataPrbsSeed: data PRBS seed. Read-write. Reset: 7FFFFh. Specifies the seed value used for creating pseudo random traffic on the data bus. This register must be written with a non-zero seed value. Recommended BIOS values for DRAM training:</p> <table border="1"> <thead> <tr> <th>CmdCount</th> <th>DataPrbsSeed</th> </tr> </thead> <tbody> <tr> <td>32</td> <td>62221h</td> </tr> <tr> <td>64</td> <td>66665h</td> </tr> <tr> <td>128</td> <td>26666h</td> </tr> <tr> <td>25644443h</td> <td></td> </tr> </tbody> </table>	CmdCount	DataPrbsSeed	32	62221h	64	66665h	128	26666h	25644443h	
CmdCount	DataPrbsSeed										
32	62221h										
64	66665h										
128	26666h										
25644443h											

D18F2x274_dct[1:0] DRAM DQ Mask Low

See D18F1x10C[DctCfgSel]. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description								
31:0	<p>DQMask[31:0]: DQ mask. Read-write. <code>DQMask[63:0] = {D18F2x278_dct[1:0][DQMask[63:32]], DQMask[31:0]}</code>. Reset: 0000_0000_0000_0000h. 1=The corresponding DQ bit will not be compared. 0=The corresponding DQ bit will be compared. See <code>D18F2x264_dct[1:0][ErrCnt]</code>.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Data[0]</td> </tr> <tr> <td>[62:1]</td> <td>Data[<code><DQMask></code>]</td> </tr> <tr> <td>[63]</td> <td>Data[63]</td> </tr> </tbody> </table>	Bit	Description	[0]	Data[0]	[62:1]	Data[<code><DQMask></code>]	[63]	Data[63]
Bit	Description								
[0]	Data[0]								
[62:1]	Data[<code><DQMask></code>]								
[63]	Data[63]								

D18F2x278_dct[1:0] DRAM DQ Mask High

Bits	Description
31:0	DQMask[63:32]: DQ mask. See: <code>D18F2x274_dct[1:0][DQMask[31:0]]</code> .

D18F2x28C_dct[1:0] DRAM Command 2

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation]. This register may only be used when D18F2x250_dct[1:0][CmdTestEnable]=1.

Bits	Description								
31	SendActCmd: send activate command. Read; write-1-only; cleared-by-hardware. 1=The DCT sends an activate command as specified by ChipSelect, Bank, and Address. This bit is cleared by hardware after the command completes.								
30	SendPchgCmd: send precharge all command. Read; write-1-only; cleared-by-hardware. The DCT sends a precharge command based on CmdAddress[10]. This bit is cleared by hardware after the command completes. 0=Command has completed. 1=If (CmdAddress[10]=1) then send a precharge all command as specified by CmdChipSelect; If (CmdAddress[10]=0) then send a precharge command as specified by CmdChipSelect, CmdBank.								
29:22	CmdChipSelect: command chip select. Read-write. Specifies the chip select. For LR-DIMMs, BIOS programs this field with the logical rank CS. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CS0</td> </tr> <tr> <td>[6:1]</td> <td>CS<CmdChipSelect></td> </tr> <tr> <td>[7]</td> <td>CS7</td> </tr> </tbody> </table>	Bit	Description	[0]	CS0	[6:1]	CS<CmdChipSelect>	[7]	CS7
Bit	Description								
[0]	CS0								
[6:1]	CS<CmdChipSelect>								
[7]	CS7								
21:19	CmdBank[2:0]: command bank [2:0]. Read-write. Specifies the bank address.								
18	Reserved.								
17:0	CmdAddress[17:0]: command address [17:0]. Read-write. Specifies the row address.								

D18F2x290_dct[1:0] DRAM Status 3

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description						
31:27	Reserved.						
26:24	ErrBeatNum: error beat number. Read-only. Indicates the data beat of the first error occurrence in the command reported by ErrCmdNum when D18F2x264_dct[1:0][ErrCnt] > 0 and D18F2x260_dct[1:0][CmdCnt] > 0. Cleared by D18F2x250_dct[1:0][ResetAllErr]. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7h-0h</td> <td><ErrBeatNum> beat</td> </tr> </tbody> </table>	Bits	Description	7h-0h	<ErrBeatNum> beat		
Bits	Description						
7h-0h	<ErrBeatNum> beat						
23:21	Reserved.						
20:0	ErrCmdNum: error command number. Read-only. Indicates the command number of the first error occurrence when D18F2x264_dct[1:0][ErrCnt] > 0 and D18F2x260_dct[1:0][CmdCnt] > 0. Cleared by D18F2x250_dct[1:0][ResetAllErr]. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>No error</td> </tr> <tr> <td>1F_FFFFh-1h</td> <td><ErrCmdNum> command</td> </tr> </tbody> </table>	Bits	Description	0	No error	1F_FFFFh-1h	<ErrCmdNum> command
Bits	Description						
0	No error						
1F_FFFFh-1h	<ErrCmdNum> command						

D18F2x294_dct[1:0] DRAM Status 4

See 2.9.1 [DCT Configuration Registers]. See 2.9.5.9.6.1 [DRAM Training Pattern Generation].

Bits	Description								
31:0	DQErr[31:0]: DQ error. Read-only. DQErr[63:0] = {D18F2x298_dct[1:0][DQErr[63:32]], DQErr[31:0]}. Reset: 0000_0000_0000_0000h. Indicates error detection status on a per bit basis when D18F2x264_dct[1:0][ErrCnt] > 0. Status is accumulated until cleared by D18F2x250_dct[1:0][ResetAllErr]. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Data[0]</td> </tr> <tr> <td>[62:1]</td> <td>Data[<DQErr>]</td> </tr> <tr> <td>[63]</td> <td>Data[63]</td> </tr> </tbody> </table>	Bit	Description	[0]	Data[0]	[62:1]	Data[<DQErr>]	[63]	Data[63]
Bit	Description								
[0]	Data[0]								
[62:1]	Data[<DQErr>]								
[63]	Data[63]								

D18F2x298_dct[1:0] DRAM Status 5

Bits	Description
31:0	DQErr[63:32]: DQ error. See: D18F2x294_dct[1:0][DQErr[31:0]].

D18F2x2E0_dct[1:0] Memory P-state Control and Status

See 2.9.1 [DCT Configuration Registers].

Bits	Description								
31	Reserved.								
30	FastMstateDis: fast M-state change disable. Read-write. Reset: 0. 1=The DCT changes MEMCLK frequency only after the NCLK frequency has changed. 0=The DCT changes MEMCLK frequency while the northbridge changes NCLK.								
29	Reserved.								
28:24	M1MemClkFreq: M1 memory clock frequency. Read-write. Reset: 00h. Specifies the frequency of the DRAM interface (MEMCLK) for memory P-state 1. See Table 132 [Memory Clock Frequency Value Definition]. The hardware enforces D18F5x84[DdrMaxRateEnf] when writes to this field occur. See D18F5x84[DdrMaxRate] and D18F5x84[DdrMaxRateEnf]. BIOS must also program D18F2x9C_x0D0F_E000_dct[1:0]_mp[1:0][Rate] for M1.								
23	Reserved.								
22:20	MxMrsEn: Mx Mrs enable. Read-write. Reset:0h. 1=The DCT writes to the DRAM MR after a memory P-state change. 0=The DCT does not write to the DRAM MR. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description, MR value</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>MR0, D18F2x2E8_dct[1:0]_mp[1:0][MxMr0]</td> </tr> <tr> <td>[1]</td> <td>MR1, D18F2x2E8_dct[1:0]_mp[1:0][MxMr1]</td> </tr> <tr> <td>[2]</td> <td>MR2, D18F2x2EC_dct[1:0]_mp[1:0][MxMr2]</td> </tr> </tbody> </table>	Bit	Description, MR value	[0]	MR0, D18F2x2E8_dct[1:0]_mp[1:0][MxMr0]	[1]	MR1, D18F2x2E8_dct[1:0]_mp[1:0][MxMr1]	[2]	MR2, D18F2x2EC_dct[1:0]_mp[1:0][MxMr2]
Bit	Description, MR value								
[0]	MR0, D18F2x2E8_dct[1:0]_mp[1:0][MxMr0]								
[1]	MR1, D18F2x2E8_dct[1:0]_mp[1:0][MxMr1]								
[2]	MR2, D18F2x2EC_dct[1:0]_mp[1:0][MxMr2]								
19:1	Reserved.								
0	CurMemPstate: current memory P-state. Read-only; updated-by-hardware. Specifies the current memory P-state. 0=M0. 1=M1.								

D18F2x2E8_dct[1:0]_mp[1:0] MRS Buffer

- See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:16	MxMr1: Mx MR1. Read-write. Reset: 0000h. Specifies the value written to DRAM MR1 after a memory P-state change. If the M1 value is the same as the M0 value, then BIOS should optimize P-state switching latency by programming <code>D18F2x2E0_dct[1:0][MxMrsEn]=0</code>
15:0	MxMr0: Mx MR0. Read-write. Reset: 0000h. Specifies the value written to DRAM MR0 after a memory P-state change. Bit 8 of MxMr0, corresponding to “DLL Reset”, should be programmed with 0b. If the M1 value is the same as the M0 value, then BIOS should optimize P-state switching latency by programming <code>D18F2x2E0_dct[1:0][MxMrsEn]=0</code>

D18F2x2EC_dct[1:0]_mp[1:0] MRS Buffer

See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:16	Reserved.
15:0	MxMr2: Mx MR2. Read-write. Reset: 0000h. Specifies the value written to DRAM MR2 after a memory P-state change. If the M1 value is the same as the M0 value, then BIOS should optimize P-state switching latency by programming <code>D18F2x2E0_dct[1:0][MxMrsEn]=0</code>

D18F2x2F0_dct[1:0]_mp[1:0] DRAM Controller Misc 3

See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:1	Reserved.
0	EffArbDis: Efficient arbitration disable. Read-write. Reset: 0. BIOS: 0. 0=The DCT optimizes the arbitration phases to improve performance under certain traffic conditions whenever the NCLK to MEMCLK ratio is less than 2:1. 1=The DCT arbitrates normally, at all NCLK:MEMCLK ratios.

D18F2x400_dct[1:0] GMC to DCT Control 0

See 2.9.1 [DCT Configuration Registers].

The GMC to DCT interface controls how DRAM bus resources are allocated and arbitrated between the MCT and the GMC. A token is the unit of available resource and is equivalent to a DCQ entry. A minimum count guarantees a number of available DCQ entries. A token limit for MCT or GMC guarantees resources are not all allocated to the GMC, or MCT respectively. Limits are configured bimodal: for normal GMC traffic and for when urgent (nominally display refresh) GMC traffic is occurring.

Bits	Description
31:16	Reserved.
15:12	Reserved.
11:8	GmcTokenLimit: GMC token limit. Read-write. Reset: 0h. BIOS: 4h. Limit of outstanding GMC tokens.
7:4	Reserved.
3:0	MctTokenLimit: MCT token limit. Read-write. Reset: 0h. BIOS: 4h. Limit of outstanding MCT tokens.

D18F2x404_dct[1:0] GMC to DCT Control 1

Reset: 0000_0000h. See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:17	Reserved.
16	UrgentTknDis: Urgent token disable. Read-write. BIOS: 0. 0=When urgent GMC traffic is requested , override the programmed values in D18F2x400_dct[1:0] and force the token scheme to heavily weight towards graphics by using the programmable token limits in D18F2x404_dct[1:0] . 1=Token scheme remains at the previously programmed non-urgent token limits in D18F2x400_dct[1:0] regardless of urgent GMC traffic.
15:12	UrGmcMinTokens: Display refresh GMC minimum tokens. Read-write. BIOS: 4h. Urgent mode minimum number of tokens assigned to the GMC.
11:8	UrGmcTokenLimit: Display refresh GMC token limit. Read-write. BIOS: 4h. Urgent mode limit of outstanding GMC tokens.
7:4	UrMctMinTokens: Display refresh MCT minimum tokens. Read-write. BIOS: 4h. Urgent mode minimum number of tokens assigned to the MCT.
3:0	UrMctTokenLimit: Display refresh MCT token limit. Read-write. BIOS: 4h. Urgent mode limit of outstanding MCT tokens. IF (!UrgentTknDis) THEN UrMctTokenLimit <= D18F2x400_dct[1:0] [MctTokenLimit].

D18F2x408_dct[1:0] GMC to DCT Control 2

See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:2	Reserved.
1	TokenAllocSelect: Token allocation select. Read-write. Reset: 0. BIOS: 0. 0=When both the MCT and GMC have less than their maximum outstanding tokens, garlic allocates tokens by alternating between each. 1= When both the MCT and GMC have less than their maximum outstanding tokens, garlic allocates tokens to whichever has less (DCQ entries + current outstanding).
0	CpuElevPrioDis: Cpu elevate priority disable. Read-write. Reset: 0. 1=Reads from MCT arbitrate with GMC traffic normally. 0=Elevate the priority of an MCT read to high. This can alleviate CPU stalls during very long graphics requests.

D18F2x420_dct[1:0] GMC to DCT FIFO Config 1

See 2.9.1 [DCT Configuration Registers].

Bits	Description
31:12	Reserved.
11:8	SbRdPtrInit: Sideband signal read pointer init value. Read-write. Reset: 4h. BIOS: 4h. Specifies the read pointer for the sideband signal FIFO between GMC and DCT.
7:4	Reserved.
3:0	CmdRdPtrInit: Command read pointer init value. Read-write. Reset: 4h. BIOS: 4h. Specifies the read pointer for command channel between GMC and DCT.

3.11 Device 18h Function 3 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space].

D18F3x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1403h.
15:0	VendorID: vendor ID. Read-only. Value: 1022h.

D18F3x04 Status/Command

Bits	Description
31:16	Status. Read-only. Reset: 0000h, except bit[20]. Bit[20] is set to indicate the existence of a PCI-defined capability block, if one exists.
15:0	Command. Read-only. Reset: 0000h.

D18F3x08 Class Code/Revision ID

Bits	Description
31:8	ClassCode. Read-only. Reset: 060000h. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only. Reset: 00h.

D18F3x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:0	HeaderTypeReg. Read-only. These bits are fixed at their default values. The header type field indicates that there are multiple functions present in this device.

D18F3x34 Capability Pointer

Bits	Description
31:8	Reserved.
7:0	CapPtr. Read-only. Value: 00h.

D18F3x40 MCA NB Control

Bits	Description
31:0	MSR0000_0410[31:0] is an alias of D18F3x40 . See MSR0000_0410[31:0] .

D18F3x44 MCA NB Configuration

See [D18F3x180 \[Extended NB MCA Configuration\]](#). It is expected that all fields of this register are programmed to the same value in all nodes, except for the fields used for link error injection: GenLinkSel, GenSubLinkSel, GenCrcErrByte1, GenCrcErrByte0.

Bits	Description
31	NbMcaLogEn: northbridge MCA log enable. Read-write. Reset: 0. 1=Enables logging (but not reporting) of NB MCA errors even if MCA is not globally enabled.
30	SyncFloodOnDramAdrParErr: sync flood on DRAM address parity error. Read-write. Reset: 0. BIOS: 1. 1=Enables sync flood on detection of a DRAM address parity error.
29	DisMstAbortCpuErrRsp: master abort CPU error response disable. Read-write. Reset: 0. 1=Disables master abort reporting through the CPU MCA error-reporting banks; Suppresses sending of RDE to CPU; Does not log any MCA information in the NB.
28	DisTgtAbortCpuErrRsp: target abort CPU error response disable. Read-write. Reset: 0. 1=Disables target abort reporting through the CPU MCA error-reporting banks; Suppresses sending of RDE to CPU; Does not log any MCA information in the NB.
27	NbMcaToMstCpuEn: machine check errors to master CPU only. Read-write. Reset: 0. BIOS: 1. 1=NB MCA errors in CMP device are only reported to the node base core (NBC); the following registers are only accessible from the NBC, non-NBC writes are ignored and reads are RAZ: MSR0000_0410 , MSR0000_0411 , MSR0000_0412 . This allows machine check handlers running on different cores to avoid coordinating accesses to the NB MCA registers. This field does not affect PCI-defined configuration space accesses to these registers, which are accessible from all cores. See 3.1 [Register Descriptions and Mnemonics] for a description of MSR space and 3 [Registers] for PCI-defined configuration space. 0=NB MCA errors may be reported to the core that originated the request, if applicable and known, and the NB MCA registers in MSR space are accessible from any core. Note: <ul style="list-style-type: none"> • When the CPU which originated the request is known, it is stored in D18F3x4C[ErrCoreId], regardless of the setting of NbMcaToMstCpuEn. See Table 224 for errors where ErrCoreId is known. • If IO originated the request, then the error is reported to the NBC, regardless of the setting of NbMcaToMstCpuEn.
26	FlagMcaCorrErr: correctable error MCA exception enable. Read-write. Reset: 0. 1=Raise a machine check exception for correctable machine check errors which are enabled in D18F3x40 .
25	DisPciCfgCpuErrRsp: PCI configuration CPU error response disable. Read-write. Reset: 0. 1=Disables generation of an error response to the core on detection of a master abort, target abort, or data error condition, and disables logging and reporting through the MCA error-reporting banks for PCI configuration accesses. For NB WDT errors on PCI configuration accesses, this prevents sending an error response to the core, but does not affect logging and reporting of the NB WDT error. See D18F3x180[DisPciCfgCpuMstAbtRsp] , which applies only to master aborts.
24	IoRdDatErrEn: IO read data error log enable. Read-write. Reset: 0. 1=Enables MCA logging and reporting of errors on transactions from IO devices upon detection of a target abort, master abort, or data error condition. 0=Errors on transactions from IO devices are not logged in MCA, although error responses to the requesting IO device may still be generated.
23:22	Reserved.

21	SyncFloodOnAnyUcErr: sync flood on any UC error. Read-write. Reset: 0. BIOS: 1. 1=Enables flooding of all links with sync packets on detection of any NB MCA error that is uncorrectable, including northbridge array errors and link protocol errors.										
20	SyncFloodOnWDT: sync flood on watchdog timer error. Read-write. Reset: 0. BIOS: 1. 1=Enables flooding of all links with sync packets on detection of a watchdog timer error.										
19:18	<p>GenSubLinkSel: sublink select for CRC error generation. Read-write. Reset: 0. Selects the sublink of a link selected by GenLinkSel to be used for CRC error injection through GenCrcErrByte0 and GenCrcErrByte1. When the link is ganged, GenSubLinkSel must be 00b. When the link is unganged, the following values indicate which sublink is selected:</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Sublink 0</td> </tr> <tr> <td>01b</td> <td>Sublink 1</td> </tr> <tr> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Sublink 0	01b	Sublink 1	10b	Reserved	11b	Reserved
<u>Bits</u>	<u>Description</u>										
00b	Sublink 0										
01b	Sublink 1										
10b	Reserved										
11b	Reserved										
17	GenCrcErrByte1: generate CRC error on byte lane 1. Read-Write. Reset: 0. 1=For ganged links (see GenSubLinkSel), a CRC error is injected on byte lane 1 of the link specified by GenLinkSel. For ganged links in retry mode or unganged links, this field is reserved, and GenCrcErrByte0 must be used. The data carried by the link is unaffected. This bit is cleared after the error has been generated.										
16	GenCrcErrByte0: generate CRC error on byte lane 0. Read-Write. Reset: 0. 1=Causes a CRC error to be injected on byte lane 0 of the link specified by GenLinkSel and the sublink specified by GenSubLinkSel. The data carried by the link is unaffected. This bit is cleared after the error has been generated.										
15:14	Reserved.										
13:12	<p>WDTBaseSel: watchdog timer time base select. Read-write. Reset: 0. Selects the time base used by the watchdog timer. The counter selected by WDTCntSel determines the maximum count value in the time base selected by WDTBaseSel.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1.31 ms</td> </tr> <tr> <td>01b</td> <td>1.28 us</td> </tr> <tr> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	1.31 ms	01b	1.28 us	10b	Reserved	11b	Reserved
<u>Bits</u>	<u>Description</u>										
00b	1.31 ms										
01b	1.28 us										
10b	Reserved										
11b	Reserved										

11:9	<p>WDTCntSel[2:0]: watchdog timer count select bits[2:0]. Read-write. Reset: 0. Selects the count used by the watchdog timer. WDTCntSel = {D18F3x180[WDTCntSel[3]], D18F3x44[WDTCntSel[2:0]]}. The counter selected by WDTCntSel determines the maximum count value in the time base selected by WDTBaseSel. WDTCntSel is encoded as:</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>4095</td> </tr> <tr> <td>0001b</td> <td>2047</td> </tr> <tr> <td>0010b</td> <td>1023</td> </tr> <tr> <td>0011b</td> <td>511</td> </tr> <tr> <td>0100b</td> <td>255</td> </tr> <tr> <td>0101b</td> <td>127</td> </tr> <tr> <td>0110b</td> <td>63</td> </tr> <tr> <td>0111b</td> <td>31</td> </tr> <tr> <td>1000b</td> <td>8191</td> </tr> <tr> <td>1001b</td> <td>16383</td> </tr> <tr> <td>1111b-1010b</td> <td>Reserved</td> </tr> </tbody> </table> <p>Because WDTCntSel is split between two registers, care must be taken when programming WDTCntSel to ensure that a reserved value is never used by the watchdog timer or undefined behavior could result.</p>	<u>Bits</u>	<u>Description</u>	0000b	4095	0001b	2047	0010b	1023	0011b	511	0100b	255	0101b	127	0110b	63	0111b	31	1000b	8191	1001b	16383	1111b-1010b	Reserved
<u>Bits</u>	<u>Description</u>																								
0000b	4095																								
0001b	2047																								
0010b	1023																								
0011b	511																								
0100b	255																								
0101b	127																								
0110b	63																								
0111b	31																								
1000b	8191																								
1001b	16383																								
1111b-1010b	Reserved																								
8	<p>WDTDis: watchdog timer disable. Read-write. Cold reset: 0. 1=Disables the watchdog timer. The watchdog timer is enabled by default and checks for NB system accesses for which a response is expected and where no response is received. If such a condition is detected the outstanding access is completed by generating an error response back to the requestor. An MCA error may also be generated if enabled in D18F3x40 [MCA NB Control].</p>																								
7	<p>IoErrDis: IO error response disable. Read-write. Reset: 0. 1=Disables setting either Error bit in link response packets to IO devices on detection of a target or master abort error condition.</p>																								
6	<p>CpuErrDis: CPU error response disable. Read-write. Reset: 0. BIOS: 1. 1=Disables generation of a read data error response to the core on detection of a target or master abort error condition.</p>																								
5	<p>IoMstAbortDis: IO master abort error response disable. Read-write. Reset: 0. 1=Signals target abort instead of master abort in link response packets to IO devices on detection of a master abort error condition. When IoMstAbortDis and D18F3x180[ChgMstAbortToNoErr] are both set, D18F3x180[ChgMstAbortToNoErr] takes precedence.</p>																								
4	<p>SyncPktPropDis: sync packet propagation disable. Read-write. Reset: 0. 1=Disables flooding of all outgoing links with sync packets when a sync packet is detected on an incoming link. Sync packets are propagated by default.</p>																								
3	<p>SyncPktGenDis: sync packet generation disable. Read-write. Reset: 0. 1=Disables flooding of all outgoing links with sync packets when a CRC error is detected on an incoming link. By default, sync packet generation for CRC errors is controlled through D18F0x84 [Link Control].</p>																								
2	<p>SyncFloodOnDramUcEcc: sync flood on uncorrectable DRAM ECC error. Read-write. Reset: 0. 1=Enables flooding of all links with sync packets on detection of an uncorrectable ECC error.</p>																								
1	<p>CpuRdDatErrEn: CPU read data error log enable. Read-write. Reset: 0. 1=Enables reporting of read data errors (master aborts and target aborts) for data destined for the CPU on this node. This bit should be clear if read data error logging is enabled for the remaining error reporting blocks in the CPU. Logging the same error in more than one block may cause a single error event to be treated as a multiple error event and cause the CPU to enter shutdown.</p>																								
0	Reserved.																								

D18F3x48 MCA NB Status Low

Bits	Description
31:0	MSR0000_0411 [31:0] is an alias of D18F3x48 . See MSR0000_0411 .

D18F3x4C MCA NB Status High

Bits	Description
31:0	MSR0000_0411 [63:32] is an alias of D18F3x4C . See MSR0000_0411 .

D18F3x50 MCA NB Address Low

Bits	Description
31:0	MSR0000_0412 [31:0] is an alias of D18F3x50 . See MSR0000_0412 [31:0].

D18F3x54 MCA NB Address High

Bits	Description
31:0	MSR0000_0412 [63:32] is an alias of D18F3x54 . See MSR0000_0412 [63:32].

D18F3x64 Hardware Thermal Control (HTC)

See 2.10.4.1 [[PROCHOT_L](#) and [Hardware Thermal Control \(HTC\)](#)]. If [D18F3xE8](#)[HtcCapable]=0 then this register is reserved.

Bits	Description				
31	Reserved.				
30:28	HtcPstateLimit: HTC P-state limit select. Read-write. Reset: Product-specific. Specifies the P-state limit of all cores when in the HTC-active state. This field uses hardware P-state numbering and is not changed on a write if the value written is greater than D18F3xDC [HwpPstateMaxVal]. No P-state limit is applied if the value written is less than D18F4x15C [NumBoostStates]. See 2.10.4.1 [PROCHOT_L and Hardware Thermal Control (HTC)] and 2.5.3.1.2.2 [Hardware P-state Numbering].				
27:24	HtcHystLmt: HTC hysteresis. Read-write. Reset: Product-specific. The processor exits the HTC-active state when Tctl is less than HTC temperature limit (HtcTmpLmt) minus HTC hysteresis (HtcHystLmt). <table border="0" style="margin-left: 20px;"> <tr> <td style="text-align: right;"><u>Bits</u></td> <td style="text-align: left;"><u>Description</u></td> </tr> <tr> <td style="text-align: right;">Fh-0h</td> <td style="text-align: left;"><HtcHystLmt*0.5></td> </tr> </table>	<u>Bits</u>	<u>Description</u>	Fh-0h	<HtcHystLmt*0.5>
<u>Bits</u>	<u>Description</u>				
Fh-0h	<HtcHystLmt*0.5>				
23	HtcSlewSel: HTC slew-controlled temperature select. Read-write. Reset: 0. 1=HTC logic is driven by the slew-controlled temperature, Tctl, specified in D18F3xA4 [Reported Temperature Control]. 0=HTC logic is driven by the measured control temperature with no slew controls.				
22:16	HtcTmpLmt: HTC temperature limit. Read-write. Reset: Product-specific. The processor enters the HTC-active state when Tctl reaches or exceeds the temperature limit defined by this register. <table border="0" style="margin-left: 20px;"> <tr> <td style="text-align: right;"><u>Bits</u></td> <td style="text-align: left;"><u>Description</u></td> </tr> <tr> <td style="text-align: right;">7Fh-00h</td> <td style="text-align: left;"><HtcTmpLmt*0.5 + 52></td> </tr> </table>	<u>Bits</u>	<u>Description</u>	7Fh-00h	<HtcTmpLmt*0.5 + 52>
<u>Bits</u>	<u>Description</u>				
7Fh-00h	<HtcTmpLmt*0.5 + 52>				

15:8	Reserved.
7	PslApicLoEn: P-state limit lower value change APIC interrupt enable. Read-write. Reset: 0. PslApicLoEn and PslApicHiEn enable interrupts using APIC330 [LVT Thermal Sensor] of each core when the active P-state limit in MSRC001_0061 [P-state Current Limit][CurPstateLimit] changes. PslApicLoEn enables the interrupt when the limit value becomes lower (indicating higher performance). PslApicHiEn enables the interrupt when the limit value becomes higher (indicating lower performance). 1=Enable interrupt.
6	PslApicHiEn: P-state limit higher value change APIC interrupt enable. Read-write. Reset: 0. See PslApicLoEn.
5	HtcActSts: HTC-active status. Read; set-by-hardware; write-1-to-clear. Reset: 0. This bit is set by hardware when the processor enters the HTC-active state.
4	HtcAct: HTC-active state. Read-only, updated-by-hardware. Reset: X. . 1=The processor is currently in the HTC-active state. 0=The processor is not in the HTC-active state.
3:1	Reserved.
0	HtcEn: HTC enable. Read-write. Reset: 0. BIOS: IF (D18F3x64[HtcTmpLmt]==0) THEN 0. ELSE 1. ENDIF. 1=HTC is enabled; the processor is capable of entering the HTC-active state.

D18F3x68 Software P-state Limit

See 2.10.4.3 [Software P-state Limit Control]. If D18F3xE8[HtcCapable]=0 then this register is reserved.

Bits	Description
31	Reserved.
30:28	SwPstateLimit: software P-state limit select. Read-write. Reset: Product-specific. Specifies a P-state limit for all cores. Uses hardware P-state numbering; see 2.5.3.1.2.2 [Hardware P-state Numbering]. Not changed on a write if the value written is greater than D18F3xDC[HwPstateMaxVal]. No P-state limit is applied if the value written is less than D18F4x15C[NumBoostStates]. See SwPstateLimitEn.
27:6	Reserved.
5	SwPstateLimitEn: software P-state limit enable. Read-write. Reset: 0. 1=SwPstateLimit is enabled.
4:0	Reserved.

D18F3x6C Data Buffer Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use D18F0x6C[RIsLnkFullTokCntImm] for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

- To ensure deadlock free operation the following minimum buffer allocations are required:
 - D18F3x6C[UpRspDBC] >= 1.
 - D18F3x6C[DnReqDBC] >= 1.
 - D18F3x6C[UpReqDBC] >= 1.
 - D18F3x6C[DnRspDBC] >= 1.
- If D18F0x84[IsocEn]=1: IsocRspDBC >= 1.
- The total number of data buffers allocated in this register and D18F3x7C must satisfy the following equation:
 - D18F3x6C[UpReqDBC] + D18F3x6C[UpRspDBC] + D18F3x6C[DnReqDBC] + D18F3x6C[DnRsp-

DBC] + D18F3x6C[IsocRspDBC] + (IF (D18F3x7C[Sri2XbarFreeRspDBC]==0) THEN (D18F3x7C[Sri2XbarFreeXreqDBC]*2) ELSE D18F3x7C[Sri2XbarFreeXreqDBC] ENDIF) + D18F3x7C[Sri2XbarFreeRspDBC] <= 16.

Bits	Description
31	Reserved.
30:28	IsocRspDBC: isochronous response data buffer count. Read-write. Cold reset: 3. BIOS: 1.
27:19	Reserved.
18:16	UpRspDBC: upstream response data buffer count. Read-write. Cold reset: 2. BIOS: 1.
15:8	Reserved.
7:6	DnRspDBC: downstream response data buffer count. Read-write. Cold reset: 2. BIOS: 1.
5:4	DnReqDBC: downstream request data buffer count. Read-write. Cold reset: 1. BIOS: 1.
3	Reserved.
2:0	UpReqDBC: upstream request data buffer count. Read-write. Cold reset: 2. BIOS: 2.

D18F3x70 SRI to XBAR Command Buffer Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use D18F0x6C[RlsLnkFullTokCntImm] for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

- To ensure deadlock free operation the following minimum buffer allocations are required:
 - D18F3x70[UpRspCBC] >= 1.
 - D18F3x70[UpPreqCBC] >= 1.
 - D18F3x70[DnPreqCBC] >= 1.
 - D18F3x70[UpReqCBC] >= 1.
 - D18F3x70[DnReqCBC] >= 1.
 - D18F3x70[DnRspCBC] >= 1.
- IF (D18F0x84[IsocEn]) THEN (D18F3x70[IsocReqCBC] >= 1).
- IF (D18F0x84[IsocEn]) THEN (D18F3x70[IsocRspCBC] >= 1).
- If D18F0x84[IsocEn]=1 and isochronous posted requests may be generated by the system: IsocPreqCBC >= 1
- The total number of SRI to XBAR command buffers allocated in this register and D18F3x7C must satisfy the following equation:
 - D18F3x70[IsocRspCBC] + D18F3x70[IsocPreqCBC] + D18F3x70[IsocReqCBC] + D18F3x70[UpRspCBC] + D18F3x70[DnPreqCBC] + D18F3x70[UpPreqCBC] + D18F3x70[DnReqCBC] + D18F3x70[DnRspCBC] + D18F3x70[UpReqCBC] + D18F3x7C[Sri2XbarFreeRspCBC] + D18F3x7C[Sri2XbarFreeXreqCBC] <= 48.

Bits	Description
31	Reserved.
30:28	IsocRspCBC: isoc response command buffer count. Read-write. Cold reset: 4. BIOS: 1.
27	Reserved.
26:24	IsocPreqCBC: isoc posted request command buffer count. Read-write. Cold reset: 1. BIOS: 0.
23	Reserved.
22:20	IsocReqCBC: isoc request command buffer count. Read-write. Cold reset: 5. BIOS: 1.

19	Reserved.
18:16	UpRspCBC: upstream response command buffer count. Read-write. Cold reset: 3. BIOS: 7.
15	Reserved.
14:12	DnPreqCBC: downstream posted request command buffer count. Read-write. Cold reset: 3. BIOS: 1.
11	Reserved.
10:8	UpPreqCBC: upstream posted request command buffer count. Read-write. Cold reset: 3. BIOS: 1.
7:6	DnRspCBC: downstream response command buffer count. Read-write. Cold reset: 2. BIOS: 1.
5:4	DnReqCBC: downstream request command buffer count. Read-write. Cold reset: 2. BIOS: 1.
3	Reserved.
2:0	UpReqCBC: upstream request command buffer count. Read-write. Cold reset: 3. BIOS: 7.

D18F3x74 XBAR to SRI Command Buffer Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use **D18F0x6C**[RIsLnkFullTokCntImm] for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

Table 178: XBAR Definitions

Term	Definition
SpqSize	Probe command queue size. SpqSize = 12.
SrqSize	SRQ (XBAR command and probe response to SRI) queue size. SrqSize = 44.
PrbRsp	SRQ entries hard allocated to probe responses. PrbRsp = 4.

- To ensure deadlock free operation the following minimum buffer allocations are required:
 - D18F3x74**[ProbeCBC] >= 2
 - D18F3x74**[UpReqCBC] >= 1
 - D18F3x74**[UpPreqCBC] >= 1
- IF (**D18F0x84**[IsocEn]) THEN (**D18F3x74**[IsocReqCBC] >= 1).
- If **D18F0x84**[IsocEn]=1 and isochronous posted requests may be generated by the system: IsocPreqCBC >= 1
- The total number of XBAR to SRI command buffers allocated in this register and **D18F3x7C** must satisfy the following equation:

$$\begin{aligned} & \mathbf{D18F3x74}[\mathbf{UpReqCBC}] + \mathbf{D18F3x74}[\mathbf{UpPreqCBC}] + \mathbf{D18F3x74}[\mathbf{DnReqCBC}] + \mathbf{D18F3x74}[\mathbf{DnPreqCBC}] + \\ & \mathbf{D18F3x74}[\mathbf{IsocReqCBC}] + \mathbf{D18F3x74}[\mathbf{IsocPreqCBC}] + \mathbf{D18F3x74}[\mathbf{DRReqCBC}] + \\ & \mathbf{D18F3x7C}[\mathbf{Xbar2SriFreeListCBC}] + \mathbf{D18F3x7C}[\mathbf{SrqExtFreeListBufCnt}] + (\mathbf{D18F3x1A0}[\mathbf{CpuCmdBufCnt}] \\ & * \mathbf{NumOfCompUnits}) + \mathbf{PrbRsp} \leq \mathbf{SrqSize} \end{aligned}$$
- The total number of SPQ (probe command) buffers allocated must satisfy the following equation:
 - (**D18F3x17C**[SPQPrbFreeCBC] + **D18F3x74**[ProbeCBC]) <= **SpqSize**.

Bits	Description
31:28	DRReqCBC: display refresh request command buffer count. Read-write. Cold reset: 0. BIOS: 0.
27	Reserved.
26:24	IsocPreqCBC: isochronous posted request command buffer count. Read-write. Cold reset: 0. BIOS: 1.

23:20	IsocReqCBC: isochronous request command buffer count. Read-write. Cold reset: 0. BIOS: 1.								
19:16	ProbeCBC: probe command buffer count. Read-write. Cold reset: Ch. BIOS: 8. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0 buffers</td> </tr> <tr> <td>Ch-1h</td> <td><ProbeCBC> buffers</td> </tr> <tr> <td>Fh-Dh</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	0h	0 buffers	Ch-1h	<ProbeCBC> buffers	Fh-Dh	Reserved.
Bits	Description								
0h	0 buffers								
Ch-1h	<ProbeCBC> buffers								
Fh-Dh	Reserved.								
15	Reserved.								
14:12	DnPreqCBC: downstream posted request command buffer count. Read-write. Cold reset: 0. BIOS: 0.								
11	Reserved.								
10:8	UpPreqCBC: upstream posted request command buffer count. Read-write. Cold reset: 1. BIOS: 1.								
7	Reserved.								
6:4	DnReqCBC: downstream request command buffer count. Read-write. Cold reset: 0. BIOS: 0.								
3	Reserved.								
2:0	UpReqCBC: upstream request command buffer count. Read-write. Cold reset: 1. BIOS: 1.								

D18F3x78 MCT to XBAR Buffer Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use [D18F0x6C\[RlsLnkFullTokCntImm\]](#) for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

- To ensure deadlock free operation the following minimum buffer allocations are required:
 ProbeCBC ≥ 1 RspCBC ≥ 1 RspDBC ≥ 2
 RspDBC \geq [D18F2x11C\[MctPrefReqLimit\]](#)+2
- To ensure deadlock free operation when online spare is enabled ([D18F2x\[5C:40\]_dct\[1:0\]\[Spare\]](#) = 1) the following minimum buffer allocation is required:
 RspCBC \geq Dh
- The total number of command buffers allocated in this register must satisfy the following equation:
[D18F3x78\[ProbeCBC\]](#) + [D18F3x78\[RspCBC\]](#) ≤ 32

Bits	Description												
31:22	Reserved.												
21:16	RspDBC: response data buffer count. Read-write. Cold reset: 20h. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>01h-00h</td> <td>Reserved</td> </tr> <tr> <td>02h</td> <td>2 Buffers</td> </tr> <tr> <td>1Fh-03h</td> <td><RspDBC> Buffers</td> </tr> <tr> <td>20h</td> <td>32 Buffers</td> </tr> <tr> <td>3Fh-21h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	01h-00h	Reserved	02h	2 Buffers	1Fh-03h	<RspDBC> Buffers	20h	32 Buffers	3Fh-21h	Reserved
Bits	Description												
01h-00h	Reserved												
02h	2 Buffers												
1Fh-03h	<RspDBC> Buffers												
20h	32 Buffers												
3Fh-21h	Reserved												
15:13	Reserved.												
12:8	ProbeCBC: probe command buffer count. Read-write. Cold reset: Ch.												
7:5	Reserved.												
4:0	RspCBC: response command buffer count. Read-write. Cold reset: 14h.												

D18F3x7C Free List Buffer Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use [D18F0x6C\[RlsLnkFullTokCntImm\]](#) for the values in the register to take effect. This is necessary even if the values are unchanged from the default values. See [D18F3x6C](#) and [D18F3x70](#).

- To ensure deadlock free operation the following minimum buffer allocations are required:
 - IF ([D18F3x7C\[Sri2XbarFreeRspCBC\]](#)==0) THEN ([D18F3x7C\[Sri2XbarFreeXreqCBC\]](#)>2).
 - IF ([D18F3x7C\[Sri2XbarFreeRspCBC\]](#)!=0) THEN ([D18F3x7C\[Sri2XbarFreeRspCBC\]](#)>2).
 - IF ([D18F3x7C\[Sri2XbarFreeRspDBC\]](#)==0) THEN ([D18F3x7C\[Sri2XbarFreeXreqDBC\]](#)>2).
 - IF ([D18F3x7C\[Sri2XbarFreeRspDBC\]](#)!=0) THEN ([D18F3x7C\[Sri2XbarFreeRspDBC\]](#)>2).
 - [D18F3x7C\[Xbar2SriFreeListCBC\]](#) >= ([D18F3x1A0\[CpuToNbFreeBufCnt\]](#) * [NumOfCompUnits](#)) + 2.

Bits	Description
31	Reserved.
30:28	Xbar2SriFreeListCBInc: XBAR to SRI free list command buffer increment. Read-write. Cold reset: 0. This field is used to add buffers to the free list pool if they are reclaimed from hard allocated entries without having to go through warm reset. This field may only be programmed after buffers have been allocated and released via D18F0x6C[RlsLnkFullTokCntImm] .
27	Reserved.
26:23	ExtSrqFreeList: extend SRQ freelist tokens. Read-write. Cold reset: 8h. BIOS: 8h. Can only be used by requests from cores to DRAM.
22:20	Sri2XbarFreeRspDBC: SRI to XBAR free response data buffer count. Read-write. Cold reset: 0. BIOS: 0.
19:16	Sri2XbarFreeXreqDBC: SRI to XBAR free request and posted request data buffer count. Read-write. Cold reset: 3h. BIOS: 5. If Sri2XbarFreeRspDBC =0h, then these buffers are shared between requests, responses and posted requests and the number of buffers allocated is two times the value of this field.
15:12	Sri2XbarFreeRspCBC: SRI to XBAR free response command buffer count. Read-write. Cold reset: Bh. BIOS: 0.
11:8	Sri2XbarFreeXreqCBC: SRI to XBAR free request and posted request command buffer count. Read-write. Cold reset: Bh. BIOS: Eh. If Sri2XbarFreeRspCBC =0h, then these buffers are shared between requests, responses and posted requests and the number of buffers allocated is two times the value of this field.
7:5	Reserved.
4:0	Xbar2SriFreeListCBC: XBAR to SRI free list command buffer count. Read-write. Cold reset: Product-specific. BIOS: 18h.

D18F3x[84:80] ACPI Power State Control

This block consists of eight identical 8-bit registers, one for each System Management Action Field (SMAF) code associated with STPCLK assertion commands from the link. Refer to the descriptions below for the associated ACPI state and system management actions for each of the 8 SMAF codes. The [SmafAct](#) fields specify the system management actions taken when the corresponding SMAF code is received. For instance, a SMAF code of 5 results in the power management actions specified by [SmafAct5](#). Some ACPI states and associated SMAF codes may not be supported in certain conditions. See [2.5 \[Power Management\]](#) for which states are supported.

When a link STPCLK assertion command is received by the processor, the power management commands specified by the register with the corresponding SMAF code are invoked. When the STPCLK deassertion command is received by the processor, the processor returns into the operational state.

In multi-node systems, these registers should be programmed identically in all nodes.

Table 179: SMAF Action Definition

Register	SmafAct	ACPI state	Description
D18F3x84[31:24]	SmafAct7	C1	Initiated when a Halt instruction is executed by processor. This does not involve the interaction with the SMC, therefore the SMC is required to never send STPCLK assertion commands with SMAF=7h.
D18F3x84[23:16]	SmafAct6	S4/S5	Initiated by a processor access to the ACPI-defined PM1_CNTa register.
D18F3x84[15:8]	SmafAct5	-	Reserved.
D18F3x84[7:0]	SmafAct4	S3	Initiated by a processor access to the ACPI-defined PM1_CNTa register.
D18F3x80[31:24]	SmafAct3	S1	Initiated by a processor access to the ACPI-defined PM1_CNTa register.
D18F3x80[23:16]	SmafAct2	-	Initiated by a processor for NB P-state changes. See 2.5.4.1 [NB P-states] .
D18F3x80[15:8]	SmafAct1	C1E, or Link init.	Initiated by an access to the ACPI-defined P_LVL3 register.
D18F3x80[7:0]	SmafAct0	C2	Initiated by a processor access to the ACPI-defined P_LVL2 register.

D18F3x80 ACPI Power State Control Low

Reset: 0000_0000h. Read-write.

Bits	Description
31:29	ClkDivisorSmafAct3 . See: ClkDivisorSmafAct0.
28:27	Reserved.
26	NbGateEnSmafAct3 . See: NbGateEnSmafAct0.
25	NbLowPwrEnSmafAct3 . See: NbLowPwrEnSmafAct0.
24	CpuPrbEnSmafAct3 . See: CpuPrbEnSmafAct0.
23:21	ClkDivisorSmafAct2 . See: ClkDivisorSmafAct0.
20:19	Reserved.
18	NbGateEnSmafAct2 . See: NbGateEnSmafAct0.
17	NbLowPwrEnSmafAct2 . See: NbLowPwrEnSmafAct0.
16	CpuPrbEnSmafAct2 . See: CpuPrbEnSmafAct0.
15:13	ClkDivisorSmafAct1 . See: ClkDivisorSmafAct0.
12:11	Reserved.
10	NbGateEnSmafAct1 . See: NbGateEnSmafAct0.
9	NbLowPwrEnSmafAct1 . See: NbLowPwrEnSmafAct0.

8	CpuPrbEnSmafAct1. See: CpuPrbEnSmafAct0.																				
7:5	<p>ClkDivisorSmafAct0: clock divisor. Read-write. Specifies the core clock frequency while in the low-power state. This divisor is relative to the current FID frequency, or:</p> <ul style="list-style-type: none"> • 100 MHz * (10h + MSRC001_00[6B:64][CpuFid]) of the current P-state specified by MSRC001_0063[CurPstate]. <p>If MSRC001_00[6B:64][CpuDid] of the current P-state indicates a divisor that is deeper than specified by this field, then no frequency change is made when entering the low-power state associated with this register.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>/1</td> <td>100b</td> <td>/16</td> </tr> <tr> <td>001b</td> <td>/2</td> <td>101b</td> <td>/128</td> </tr> <tr> <td>010b</td> <td>/4</td> <td>110b</td> <td>/512</td> </tr> <tr> <td>011b</td> <td>/8</td> <td>111b</td> <td>Turn off clocks</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	/1	100b	/16	001b	/2	101b	/128	010b	/4	110b	/512	011b	/8	111b	Turn off clocks
Bits	Description	Bits	Description																		
000b	/1	100b	/16																		
001b	/2	101b	/128																		
010b	/4	110b	/512																		
011b	/8	111b	Turn off clocks																		
4:3	Reserved.																				
2	NbGateEnSmafAct0: northbridge gate enable. Read-write. This bit does not control hardware. NbLowPwrEn is required to be set if this bit is set.																				
1	NbLowPwrEnSmafAct0: Northbridge low-power enable. Read-write. 1=The NB clock is ramped down to the divisor specified by D18F3xD4[NbClkDiv] and DRAM is placed into self-refresh mode when LDTSTOP_L is asserted while in the low-power state.																				
0	<p>CpuPrbEnSmafAct0: CPU direct probe enable. Read-write. Specifies how probes are handled while in the low-power state. 0=When the probe request comes into the NB, the core clock is brought up to the COF (based on the current P-state), all outstanding probes are completed, the core waits for a hysteresis time based on D18F3xD4[ClkRampHystSel], and then the core clock is brought down to the frequency specified by ClkDivisor. 1=The core clock does not change frequency; the probe is handled at the frequency specified by ClkDivisor; this may only be set if:</p> <ul style="list-style-type: none"> • ClkDivisor specifies a divide-by 1, 2, 4, 8, or 16 and NbCof <= 3.2 GHz • ClkDivisor specifies a divide-by 1, 2, 4, or 8 and NbCof >= 3.4 GHz <p>This bit also specifies functionality of the timer used for cache flushing during halt. See D18F3xDC[CacheFlushOnHaltTmr].</p> <ul style="list-style-type: none"> • If D18F3x[84:80][CpuPrbEnSmafAct7]=0 and D18F3xDC[IgnCpuPrbEn]=0, only the time when the core is halted and has its clocks ramped up to service probes is counted. • If D18F3x[84:80][CpuPrbEnSmafAct7]=1 or D18F3xDC[IgnCpuPrbEn]=1, all of the time the core is halted is counted. 																				

D18F3x84 ACPI Power State Control High

Reset: 0000_0000h. Read-write.

Bits	Description
31:29	ClkDivisorSmafAct7. See: D18F3x80[ClkDivisorSmafAct0].
28:27	Reserved.
26	NbGateEnSmafAct7. See: D18F3x80[NbGateEnSmafAct0].
25	NbLowPwrEnSmafAct7. See: D18F3x80[NbLowPwrEnSmafAct0].
24	CpuPrbEnSmafAct7. See: D18F3x80[CpuPrbEnSmafAct0].
23:21	ClkDivisorSmafAct6. See: D18F3x80[ClkDivisorSmafAct0].
20:19	Reserved.
18	NbGateEnSmafAct6. See: D18F3x80[NbGateEnSmafAct0].

17	NbLowPwrEnSmafAct6. See: D18F3x80[NbLowPwrEnSmafAct0] .
16	CpuPrbEnSmafAct6. See: D18F3x80[CpuPrbEnSmafAct0] .
15:13	ClkDivisorSmafAct5. See: D18F3x80[ClkDivisorSmafAct0] .
12:11	Reserved.
10	NbGateEnSmafAct5. See: D18F3x80[NbGateEnSmafAct0] .
9	NbLowPwrEnSmafAct5. See: D18F3x80[NbLowPwrEnSmafAct0] .
8	CpuPrbEnSmafAct5. See: D18F3x80[CpuPrbEnSmafAct0] .
7:5	ClkDivisorSmafAct4. See: D18F3x80[ClkDivisorSmafAct0] . BIOS: 111b.
4:3	Reserved.
2	NbGateEnSmafAct4. See: D18F3x80[NbGateEnSmafAct0] .
1	NbLowPwrEnSmafAct4. See: D18F3x80[NbLowPwrEnSmafAct0] . BIOS: 1.
0	CpuPrbEnSmafAct4. See: D18F3x80[CpuPrbEnSmafAct0] .

D18F3x88 NB Configuration Low

Bits	Description
31:0	MSRC001_001F [31:0] is an alias of D18F3x88 . See MSRC001_001F .

D18F3x8C NB Configuration High

Bits	Description
31:0	MSRC001_001F [63:32] is an alias of D18F3x8C . See MSRC001_001F .

D18F3xA0 Power Control Miscellaneous

Bits	Description
31	CofVidProg: COF and VID of P-states programmed. Read-only. Reset: Product-specific. 1=Out of cold reset, the VID, FID, and DID values of the P-state registers specified by MSRC001_0071 [StartupPstate] and D18F5x174 [StartupNbPstate] have been applied to the processor. 0=Out of cold reset, the boot VID is applied to all processor power planes, the NB clock plane is set to 800 MHz (with a FID of 04h=800 MHz and a DID of 0b) and core CPU clock planes are set to 800 MHz (with a FID of 00h=1.6 GHz and a DID of 1h). Registers containing P-state information such as FID, DID, and VID values are valid out of cold reset independent of the state of D18F3xA0 [CofVidProg]. BIOS must transition the processor to a valid P-state out of cold reset when D18F3xA0 [CofVidProg]=0. See 2.5.3.1.7 [BIOS Requirements for Core P-state Initialization and Transitions].
30:28	Reserved.
27:16	ConfigId: Configuration identifier. Read-only. Reset: Product-specific. Specifies the configuration ID associated with the product.
15	Reserved.

14	Svi2HighFreqSel: SVI high frequency select. Read-write. Cold reset: 0. BIOS: 1. 0=3.4 MHz. 1=20 MHz. Writes to this field take effect at the next SVI command boundary. If 20 MHz is supported by the VRM, BIOS should program this to 1 prior to any VID transitions. Once this bit is set, it should not be cleared until the next cold reset. See 2.5.1.1 [Serial VID Interface].																				
13:11	<p>PLLLockTime: PLL synchronization lock time. Read-write. Reset: 0. BIOS: 001b. If a P-state change occurs that applies a new FID to the PLL, this field specifies the time required for the PLL to lock to the new frequency.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>1 us</td> <td>100b</td> <td>8 us</td> </tr> <tr> <td>001b</td> <td>2 us</td> <td>101b</td> <td>16 us</td> </tr> <tr> <td>010b</td> <td>3 us</td> <td>110b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>4 us</td> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	1 us	100b	8 us	001b	2 us	101b	16 us	010b	3 us	110b	Reserved	011b	4 us	111b	Reserved
Bits	Description	Bits	Description																		
000b	1 us	100b	8 us																		
001b	2 us	101b	16 us																		
010b	3 us	110b	Reserved																		
011b	4 us	111b	Reserved																		
10:9	Reserved.																				
8	PsiVid[7]. Read-write. Reset: 0. BIOS: 2.5.1.3.1.1. See PsiVid[6:0].																				
7	PsiVidEn: PSI_L VID enable. Read-write. Reset: 0. BIOS: 2.5.1.3.1.1. This bit specifies how PSI_L is controlled. This signal may be used by the voltage regulator to improve efficiency while in reduced power states. 1=Control over the PSI_L signal is as specified by the PsiVid field of this register. 0=PSI_L is always high. See 2.5.1.3.1 [PSIx_L Bit].																				
6:0	PsiVid[6:0]: PSI_L VID threshold. Read-write. Reset: 0. BIOS: 2.5.1.3.1.1. PsiVid[7:0] = {PsiVid[7], PsiVid[6:0]}. When enabled by PsiVidEn, PsiVid[7:0] specifies the threshold value of the VID code generated by the processor, which in turn determines the state of PSI0_L. When the VID code generated by the processor is less than PsiVid[7:0] (i.e., the VID code is specifying a higher voltage level than the PsiVid-specified voltage level), then PSI0_L is high; when the VID code is greater than or equal to PsiVid[7:0], PSI0_L is driven low. See 2.5.1.3.1 [PSIx_L Bit].																				

D18F3xA4 Reported Temperature Control

See 2.10.1 [The Tctl Temperature Scale] and 2.10.2 [Temperature Slew Rate Control].

Bits	Description																				
31:21	<p>CurTmp: current temperature. IF (D18F3xA4[CurTmpTjSel]==11b) THEN Read-write. ELSE Read-only, updated-by-hardware. ENDIF. Reset: X. Provides the current control temperature, Tctl, after the slew-rate controls have been applied.</p> <p>IF (D18F3xA4[CurTmpTjSel]!=11b) THEN</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000h</td> <td>0</td> </tr> <tr> <td>001h</td> <td>0.125</td> </tr> <tr> <td>7FEh-002h</td> <td><CurTmp*0.125></td> </tr> <tr> <td>7FFh</td> <td>255.875</td> </tr> </tbody> </table> <p>ELSE</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000h</td> <td>-49</td> </tr> <tr> <td>001h</td> <td>-48.875</td> </tr> <tr> <td>7FEh-002h</td> <td><(CurTmp*0.125)-49></td> </tr> <tr> <td>7FFh</td> <td>206.875</td> </tr> </tbody> </table> <p>ENDIF.</p>	Bits	Description	000h	0	001h	0.125	7FEh-002h	<CurTmp*0.125>	7FFh	255.875	Bits	Description	000h	-49	001h	-48.875	7FEh-002h	<(CurTmp*0.125)-49>	7FFh	206.875
Bits	Description																				
000h	0																				
001h	0.125																				
7FEh-002h	<CurTmp*0.125>																				
7FFh	255.875																				
Bits	Description																				
000h	-49																				
001h	-48.875																				
7FEh-002h	<(CurTmp*0.125)-49>																				
7FFh	206.875																				

20	TcenPwrDnCc6En: TCEN power down CC6 Enable. Read-write. Cold reset: 0. 1=A TCEN remains powered down if the associated compute-unit is in CC6 when the TCEN power down delay timer expires. 0=TCEN powers up immediately after TCEN power down delay timer expires. Field has no effect if CC6-based TCEN power down feature is disabled by hardware.										
19:18	Reserved.										
17:16	CurTmpTjSel: Current temperature select. Read-write. Reset: 00. These bits may be used for diagnostic software. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>CurTmp provides the read-only Tctl value.</td> </tr> <tr> <td>01b</td> <td>Reserved.</td> </tr> <tr> <td>10b</td> <td>Reserved.</td> </tr> <tr> <td>11b</td> <td>CurTmp is a read-write register that specifies a value used to create Tctl. The two LSB's are read-only zero.</td> </tr> </tbody> </table>	Bits	Description	00b	CurTmp provides the read-only Tctl value.	01b	Reserved.	10b	Reserved.	11b	CurTmp is a read-write register that specifies a value used to create Tctl. The two LSB's are read-only zero.
Bits	Description										
00b	CurTmp provides the read-only Tctl value.										
01b	Reserved.										
10b	Reserved.										
11b	CurTmp is a read-write register that specifies a value used to create Tctl. The two LSB's are read-only zero.										
15:13	Reserved.										
12:8	PerStepTimeDn: per step time down. Read-write. Cold reset: 18h. BIOS: 0Fh. Specifies the time that measured temperature must remain below Tctl before applying a 0.125 downward step. See PerStepTimeUp for encodings.										
7	TmpSlewDnEn: temperature slew downward enable. Read-write. Cold reset: 0. BIOS: 1. 1=Downward slewing enabled. 0=Downward slewing disabled.										
6:5	TmpMaxDiffUp: temperature maximum difference up. Read-write. Cold reset: 00b. BIOS: 11b. Specifies the maximum difference, (measured temperature - Tctl), when Tctl immediately updates to the measured temperature. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>0.0 (disable upward slew)</td> </tr> <tr> <td>01b</td> <td>1.0</td> </tr> <tr> <td>10b</td> <td>3.0</td> </tr> <tr> <td>11b</td> <td>9.0</td> </tr> </tbody> </table>	Bits	Description	00b	0.0 (disable upward slew)	01b	1.0	10b	3.0	11b	9.0
Bits	Description										
00b	0.0 (disable upward slew)										
01b	1.0										
10b	3.0										
11b	9.0										
4:0	PerStepTimeUp: per 1/8th degree step time up. Read-write. Cold reset: 00h. BIOS: 0Fh. Specifies the time that measured temperature must remain above Tctl before applying a 0.125 upward step.. It is encoded as follows: <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>1Fh-00h</td> <td>$\langle (\text{PerStepTimeUp}[2:0] + 1) * 10^{\text{PerStepTimeUp}[4:3]} \rangle$ ms, ranging from 1 ms to 8000 ms.</td> </tr> </tbody> </table>	Bits	Definition	1Fh-00h	$\langle (\text{PerStepTimeUp}[2:0] + 1) * 10^{\text{PerStepTimeUp}[4:3]} \rangle$ ms, ranging from 1 ms to 8000 ms.						
Bits	Definition										
1Fh-00h	$\langle (\text{PerStepTimeUp}[2:0] + 1) * 10^{\text{PerStepTimeUp}[4:3]} \rangle$ ms, ranging from 1 ms to 8000 ms.										

D18F3xA8 Pop Up and Down P-states

Bits	Description
31:29	PopDownPstate. Read-write. Reset: D18F3xDC[HwPstateMaxVal]. BIOS: D18F3xDC[HwPstateMaxVal]. Specifies the pop-down P-state number. This field uses hardware P-state numbering. See 2.5.3.2.3.3 [Core C6 (CC6) State].
28:0	Reserved.

D18F3xD4 Clock Power/Timing Control 0

Bits	Description																																				
31	NbClkDivApplyAll. Read-write. Cold reset: 0. BIOS: 1. See NbClkDiv.																																				
30:28	<p>NbClkDiv: NB clock divisor. Read-write. Cold reset: Product-specific. BIOS: 100b. Specifies the NB CLK divisor associated with D18F3x[84:80] [ACPI Power State Control][NbLowPwrEn]. This divisor is applied while LDTSTOP is asserted if the corresponding core CLK divisor, D18F3x[84:80][ClkDivisor], is set to “turn off clocks” or if NbClkDivApplyAll=1; otherwise, the divisor specified by D18F3x[84:80][ClkDivisor] is applied. This divisor is relative to the current NB FID frequency, or:</p> <ul style="list-style-type: none"> 100 MHz * (4 + D18F5x1[6C:60][NbFid]). <p>If D18F5x1[6C:60][NbDid] of the current P-state indicates a divisor that is lower than specified by this field, then no NB frequency change is made when entering the low-power state associated with this register (i.e., if this field specifies a divide-by 1 and the DID is divide-by 2, then the divisor remains 2 while in the low-power state). This field is encoded as follows:</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Divide-by 1</td> <td>100b</td> <td>Divide-by 16</td> </tr> <tr> <td>001b</td> <td>Divide-by 2</td> <td>101b</td> <td>Reserved</td> </tr> <tr> <td>010b</td> <td>Divide-by 4</td> <td>110b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>Divide-by 8</td> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	Divide-by 1	100b	Divide-by 16	001b	Divide-by 2	101b	Reserved	010b	Divide-by 4	110b	Reserved	011b	Divide-by 8	111b	Reserved																
Bits	Description	Bits	Description																																		
000b	Divide-by 1	100b	Divide-by 16																																		
001b	Divide-by 2	101b	Reserved																																		
010b	Divide-by 4	110b	Reserved																																		
011b	Divide-by 8	111b	Reserved																																		
27:24	<p>PowerStepUp. Read-write. Cold reset: 0000b. BIOS: 1000b. Specifies the rate at which blocks of compute unit and NB logic are gated on while the processor transitions from a quiescent state to an active state as part of a power management state transition. There are about 15 steps in this transition for each compute unit and about 5 steps for the NB for the PowerStepDown and PowerStepUp transitions. So the total transition time for a single compute unit is about 15 times the time specified by PowerStepDown and PowerStepUp and the transition time for the NB is about 5 times the time specified by PowerStepDown and PowerStepUp. Use of longer transition times may help reduce voltage transients associated with power state transitions. The bits for PowerStepUp and PowerStepDown are encoded as follows:</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>Reserved.</td> <td>1000b</td> <td>50 ns</td> </tr> <tr> <td>0001b</td> <td>Reserved.</td> <td>1001b</td> <td>Reserved.</td> </tr> <tr> <td>0010b</td> <td>Reserved.</td> <td>1010b</td> <td>Reserved.</td> </tr> <tr> <td>0011b</td> <td>100 ns</td> <td>1011b</td> <td>Reserved.</td> </tr> <tr> <td>0100b</td> <td>90 ns</td> <td>1100b</td> <td>Reserved.</td> </tr> <tr> <td>0101b</td> <td>80 ns</td> <td>1101b</td> <td>Reserved.</td> </tr> <tr> <td>0110b</td> <td>70 ns</td> <td>1110b</td> <td>Reserved.</td> </tr> <tr> <td>0111b</td> <td>60 ns</td> <td>1111b</td> <td>Reserved.</td> </tr> </tbody> </table> <ul style="list-style-type: none"> If PowerStepDown or PowerStepUp are programmed to greater than 50 ns, then the value applied to the NB is clipped to 50 ns. The compute unit steps are not clipped. 	Bits	Description	Bits	Description	0000b	Reserved.	1000b	50 ns	0001b	Reserved.	1001b	Reserved.	0010b	Reserved.	1010b	Reserved.	0011b	100 ns	1011b	Reserved.	0100b	90 ns	1100b	Reserved.	0101b	80 ns	1101b	Reserved.	0110b	70 ns	1110b	Reserved.	0111b	60 ns	1111b	Reserved.
Bits	Description	Bits	Description																																		
0000b	Reserved.	1000b	50 ns																																		
0001b	Reserved.	1001b	Reserved.																																		
0010b	Reserved.	1010b	Reserved.																																		
0011b	100 ns	1011b	Reserved.																																		
0100b	90 ns	1100b	Reserved.																																		
0101b	80 ns	1101b	Reserved.																																		
0110b	70 ns	1110b	Reserved.																																		
0111b	60 ns	1111b	Reserved.																																		
23:20	PowerStepDown. Read-write. Cold reset: 0000b. BIOS: 1000b. This specifies the rate at which blocks of compute unit and NB logic are gated off while the processor transitions from an active state to a quiescent state as part of a power management state transition.																																				
19:15	Reserved.																																				
14	<p>CacheFlushImmOnAllHalt: cache flush immediate on all halt. Read-write. Cold reset: 0. BIOS: 0. 1=Flush the caches immediately when all cores in a package have halted. The following condition must be true in order for the caches to be flushed:</p> <ul style="list-style-type: none"> D18F4x11[C:8][CacheFlushEn]=1 for the corresponding C-state action field on all cores. 																																				

13	Reserved.
12	ClkRampHystCtl: clock ramp hysteresis control. Read-write. Cold reset: 0. Specifies the time base for ClkRampHystSel when D18F3x[84:80][CpuPrbEn]=0 . 0=320 ns. 1=1.28 us.
11:8	<p>ClkRampHystSel: clock ramp hysteresis select. Read-write. Cold reset: 0h. BIOS: Fh. When the core(s) are in the stop-grant or halt state and a probe request is received, the core clock may need to be brought up to service the probe.</p> <ul style="list-style-type: none"> • If D18F3x[84:80][CpuPrbEn]=0 for the low-power state, then this field specifies how long the core clock is left up to service additional probes before being brought back down. Each time a probe request is received, the hysteresis timer is reset such that the period of time specified by this field must expire with no probe request before the core clock is brought back down. The hysteresis time is encoded as (the time base specified by D18F3xD4[ClkRampHystCtl]) * (1 + ClkRampHystSel). • If D18F3x[84:80][CpuPrbEn]=1 for the low-power state, and for requests to change core P-states, then this field specifies a fixed amount of time to allow for probes to be serviced after completing the transition of each core. If, for example, two cores enter stop-grant or halt at the same time, then (1) the first core would complete the transition to the low power state, (2) probe traffic would be serviced for the time specified by this field, (3) the second core would complete the transition to the low power state, and (4) probe traffic would be serviced for the time specified by this field (and afterwards, until the next power state transition). For this purpose, values range from 0h=40 ns to Fh=640 ns, encoded as 40 ns * (1 + ClkRampHystSel).
7:6	Reserved.
5:0	MaxSwPstateCpuCof: maximum software P-state core COF. Read-only. Cold reset: Product-specific. Specifies the maximum CPU COF supported by the processor in a software P-state. The maximum frequency is 100 MHz * MaxSwPstateCpuCof, if MaxSwPstateCpuCof is greater than zero; if MaxSwPstateCpuCof = 00h, then there is no frequency limit. Any attempt to change a software P-state CPU COF to a frequency greater than specified by this field is ignored. See 2.5.3.1.2.1 [Software P-state Numbering] .

D18F3xD8 Clock Power/Timing Control 1

See [2.5.1.4 \[Voltage Transitions\]](#).

Bits	Description																		
31:7	Reserved.																		
6:4	<p>VSRampSlamTime. Read-write. Cold reset: 000b. BIOS: Voltage Ramp Time¹. Specifies the time the processor waits for voltage transitions to complete before beginning an additional voltage change or a frequency change.</p> <p>Wait time = (VSRampSlamTime / 12.5mV) * ABS(destination voltage - current voltage). Wait time = (VSRampSlamTime / 15mV) * ABS(destination voltage - current voltage).</p> <table border="1"> <thead> <tr> <th>tionBits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>5.00 us</td> </tr> <tr> <td>001b</td> <td>3.75 us</td> </tr> <tr> <td>010b</td> <td>3.00 us</td> </tr> <tr> <td>011b</td> <td>2.40 us</td> </tr> <tr> <td>100b</td> <td>2.00 us</td> </tr> <tr> <td>101b</td> <td>1.50 us</td> </tr> <tr> <td>110b</td> <td>1.20 us</td> </tr> <tr> <td>111b</td> <td>1.00 us</td> </tr> </tbody> </table> <p>1. Voltage Ramp Time = The maximum time to change VDD or VDDNB 15mV rounded to the next higher encoding.</p>	tionBits	Description	000b	5.00 us	001b	3.75 us	010b	3.00 us	011b	2.40 us	100b	2.00 us	101b	1.50 us	110b	1.20 us	111b	1.00 us
tionBits	Description																		
000b	5.00 us																		
001b	3.75 us																		
010b	3.00 us																		
011b	2.40 us																		
100b	2.00 us																		
101b	1.50 us																		
110b	1.20 us																		
111b	1.00 us																		
3:0	Reserved.																		

D18F3xDC Clock Power/Timing Control 2

Bits	Description																		
31:27	Reserved.																		
26	IgnCpuPrbEn: ignore CPU probe enable. Read-write. Cold reset: 0. BIOS: 0. See D18F3x[84:80][CpuPrbEn] and D18F4x11[C:8][CpuPrbEn] .																		
25:19	CacheFlushOnHaltTmr: cache flush on halt timer. Read-write. Cold reset: 00h. BIOS: 14h. Specifies how long each core needs to stay in a C-state before it flushes its caches. See CacheFlushOnHaltCtl , D18F3x[84:80][CpuPrbEn] , and D18F4x11[C:8][CacheFlushTmrSel] . <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>5.12 us</td> </tr> <tr> <td>7Fh-01h</td> <td>(<CacheFlushOnHaltTmr> * 10.24us) - 5.12us <= Time <= <CacheFlushOnHaltTmr> * 10.24 us</td> </tr> </tbody> </table>	Bits	Description	00h	5.12 us	7Fh-01h	(<CacheFlushOnHaltTmr> * 10.24us) - 5.12us <= Time <= <CacheFlushOnHaltTmr> * 10.24 us												
Bits	Description																		
00h	5.12 us																		
7Fh-01h	(<CacheFlushOnHaltTmr> * 10.24us) - 5.12us <= Time <= <CacheFlushOnHaltTmr> * 10.24 us																		
18:16	CacheFlushOnHaltCtl: cache flush on halt control. Read-write. Cold reset: 000b. BIOS: 111b. If CacheFlushOnHaltCtl != 0, then cache flush on halt is enabled. CacheFlushOnHaltCtl also specifies what core clock divisor is used after the caches have been flushed. See D18F4x11[C:8][CacheFlushTmrSelCstAct0] . <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>/1.</td> </tr> <tr> <td>001b</td> <td>/2</td> </tr> <tr> <td>010b</td> <td>/4</td> </tr> <tr> <td>011b</td> <td>/8</td> </tr> <tr> <td>100b</td> <td>/16</td> </tr> <tr> <td>101b</td> <td>Reserved</td> </tr> <tr> <td>110b</td> <td>Reserved</td> </tr> <tr> <td>111b</td> <td>Turn off clocks</td> </tr> </tbody> </table> See D18F3x[84:80] and D18F4x11[C:8] for clock divisor specifications that are in effect during a C-state before the caches have been flushed. See 2.5.3.2.3.1 [C-state Probes and Cache Flushing] .	Bits	Description	000b	/1.	001b	/2	010b	/4	011b	/8	100b	/16	101b	Reserved	110b	Reserved	111b	Turn off clocks
Bits	Description																		
000b	/1.																		
001b	/2																		
010b	/4																		
011b	/8																		
100b	/16																		
101b	Reserved																		
110b	Reserved																		
111b	Turn off clocks																		
15	Reserved.																		
14:12	NbsynPtrAdj: NB/core synchronization FIFO pointer adjust. Read-write. Cold reset: 000b. BIOS: 101b. Changes to this field take effect after any of the following events: <ul style="list-style-type: none"> • Warm reset. • All compute units perform a P-state transition. • An NB P-state transition. There is a synchronization FIFO between the NB clock domain and core clock domains. At cold reset, the read pointer and write pointer for each of these FIFOs is positioned conservatively, such that FIFO latency may be greater than is necessary. This field may be used to position the read pointer and write pointer of each FIFO closer to each other such that latency is reduced. Each increment of this field represents one clock cycle of whichever is the slower clock (longer period) between the NB clock and the core clock. Values less than the recommended value are allowed; values greater than the recommended value are illegal. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>5h-0h</td> <td>Position the read pointer <NbsynPtrAdj> clock cycles closer to the write pointer</td> </tr> <tr> <td>7h-6h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	5h-0h	Position the read pointer <NbsynPtrAdj> clock cycles closer to the write pointer	7h-6h	Reserved												
Bits	Description																		
5h-0h	Position the read pointer <NbsynPtrAdj> clock cycles closer to the write pointer																		
7h-6h	Reserved																		
11	Reserved.																		

10:8	HwPstateMaxVal: P-state maximum value. Read-write. IF ((D18F3xE8[HtcCapable]==1) && (D18F3x64[HtcTmpLmt]!=0) && (D18F3x64[HtcPstateLimit] > HwPstateMaxVal)) THEN BIOS: D18F3x64[HtcPstateLimit]. ENDIF. Cold reset: specified by the reset state of MSRC001_00[6B:64][PstateEn]; the cold reset value is the highest P-state number corresponding to the MSR in which PstateEn is set (e.g., if MSRC001_0064 and MSRC001_0065 have this bit set and the others do not, then PstateMaxVal=1; if MSRC001_0064 has this bit set and the others do not, then PstateMaxVal=0). This specifies the highest P-state value (lowest performance state) supported by the hardware. This field must not be written to a value less (higher performance) than MSRC001_0071[CurPstateLimit]. See MSRC001_0061[PstateMaxVal]. This field uses hardware P-state numbering. See 2.5.3.1.2.2 [Hardware P-state Numbering].
7:0	Reserved.

D18F3xE4 Thermtrip Status

Bits	Description
31	SwThermtp: software THERMTRIP. Write-1-only; cleared-by-hardware. Reset: 0. Writing a 1 to this bit position induces a THERMTRIP event. This bit returns 0 when read. This is a diagnostic bit, and it should be used for testing purposes only.
30:6	Reserved.
5	ThermtpEn: THERMTRIP enable. Read-only. Reset: Product-specific. 1=The THERMTRIP state is supported. See 2.10.4.4 [THERMTRIP].
4	Reserved.
3	ThermtpSense: THERMTRIP sense. Read-only. Cold reset: 0. 1=The processor temperature exceeded the THERMTRIP value (regardless as to whether the THERMTRIP state is enabled). This bit is also set when the diagnostic bit SwThermtp = 1.
2	Reserved.
1	Thermtp: THERMTRIP. Read-only. Cold reset: 0. 1=The processor has entered the THERMTRIP state.
0	Reserved.

D18F3xE8 Northbridge Capabilities

Read-only. Value: Product-specific. Unless otherwise specified, 1=The feature is supported by the processor; 0=The feature is not supported.

Bits	Description
31:29	Reserved.
28:26	Reserved.
25	L3Capable. 1=Specifies that an L3 cache is present. See CPUID Fn8000_0006_EDX.
24	MemPstateCap: memory P-state capable.
23:20	Reserved.
19	x2Apic: x2APIC capability.
18:16	Reserved.
15	Reserved.

14	MultVidPlane: multiple VID plane capable. Value: 1.
13:12	Reserved.
11	Reserved.
10	HtcCapable: HTC capable. This affects D18F3x64 and D18F3x68 .
9	SvmCapable: SVM capable.
8	MctCap: memory controller (on the processor) capable. Value: 1.
7:5	Reserved.
4	ChipKill: chipkill ECC capable.
3	ECC: ECC capable.
2	EightNode: Eight-node multi-processor capable.
1	DualNode: Dual-node multi-processor capable.
0	Reserved.

D18F3xF0 DEV Capability Header Register

The DEV secure loader function is configured through [D18F3xF4](#) and [D18F3xF8](#). The register number (i.e., the number that follows F8_x in the register mnemonic) is specified by [D18F3xF4\[DevFunction\]](#). Access to this register is accomplished as follows:

- Reads: Write the register number to [D18F3xF4\[DevFunction\]](#). Read the register contents from [D18F3xF8](#).
- Writes: Write the register number to [D18F3xF4\[DevFunction\]](#). Write the register contents to [D18F3xF8](#).

IF ([D18F3xE8\[SvmCapable\]](#)==0) THEN

Bits	Description
31:0	Reserved.

ELSE

Bits	Description
31:19	Reserved. Reset: 0002h.
18:16	CapType: DEV capability block type. Read-only. Reset: 000b. Specifies the layout of the Capability Block.
15:8	CapPtr: capability pointer. Read-only. Reset: 00h. Indicates that this is the last capability block.
7:0	CapId: capability ID. Read-only. Reset: 0Fh. Indicates a DEV capability block.

ENDIF.

D18F3xF4 DEV Function Register

Reset: 0000_0000h.

IF ([D18F3xE8\[SvmCapable\]](#)==0) THEN

Bits	Description
31:0	Reserved.

ELSE

Bits	Description
31:16	Reserved.
15:8	DevFunction. Read-write. See D18F3xF0 . The valid value for this field is 4h. Writing invalid values may result in undefined behavior.
7:0	Reserved.

ENDIF.

D18F3xF8 DEV Data Port

Reset: 0000_0000h. See [D18F3xF0](#) for details about this port.

IF ([D18F3xE8](#)[SvmCapable]==0) THEN

Bits	Description
31:0	Reserved.

ELSE

Bits	Description
31:0	DevData. Read-write. See D18F3xF8_x4 .

ENDIF.

D18F3xF8_x4 DEV Secure Loader Control Register

Reset: 0000_0000h.

Bits	Description
31:9	Reserved.
8	SecureGfxMode: secure graphics mode. RAZ; write-0-only. Secure graphics mode is entered when an SKINIT instruction is executed. In this mode, all accesses to memory except for accesses to the frame buffer from the GPU can be checked by the IOMMU if IOMMU is enabled. This bit always reads as a zero so software should check if SDev or IoDis is set to determine whether the processor is in secure graphics mode.
7:6	Reserved.
5	SDev: secure loader DEV protection enable. Read; write-0-only; set-by-hardware. This bit is set when an SKINIT instruction is executed. 1=The memory region associated with the SKINIT instruction is protected from DMA access.
4:3	Reserved.
2	IoDis: upstream IO disable. Read-write; set-by-hardware. This bit is set when an SKINIT instruction is executed. 1=Upstream IO-space accesses are regarded as DEV protection violations.
1:0	Reserved.

D18F3xFC CPUID Family/Model

[CPUID Fn0000_0001_EAX](#), [CPUID Fn8000_0001_EAX](#) are an alias of [D18F3xFC](#).

Bits	Description
31:28	Reserved.
27:20	ExtFamily: extended family. Read-only. Reset: 06h.
19:16	ExtModel: extended model. Read-only. Value: 1h.
15:12	Reserved.
11:8	BaseFamily. Read-only. Reset: Fh.
7:4	BaseModel. Read-only. Value: Product-specific.
3:0	Stepping. Read-only. Value: Product-specific.

D18F3x140 SRI to XCS Token Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use [D18F0x6C\[RlsLnkFullTokCntImm\]](#) for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

[D18F3x140](#), [D18F3x144](#), and [D18F3x148](#) specify the number of XCS (XBAR command scheduler) entries assigned to each virtual channel within each source port. See [2.8 \[Northbridge \(NB\)\]](#). The totals of SRI, MCT and the links must not exceed the number of XCS entries. The default totals are:

- SRI: 22
- MCT: 10
- Upstream channel: 8.
- Total: 40, which is the total number of entries supported by XCS.

The defaults for [D18F3x140](#), [D18F3x148](#), do not allocate any tokens in the isochronous channel. If isochronous flow control mode (IFCM) is enabled ([D18F0x84\[IsocEn\]](#)), then the XCS token counts must be changed.

- If IFCM is enabled, then [D18F3x140\[IsocReqTok and IsocRspTok\]](#) must each be non-zero. If isochronous posted requests may be generated in the system, then [D18F3x140\[IsocPreqTok\]](#) must also be non-zero.
- If an IOMMU is present, [D18F3x148\[IsocReqTok\]](#) must be non-zero.

Bits	Description
31:24	Reserved.
23:20	FreeTok: free tokens. Read-write. Cold reset: Ah. BIOS: Ah. The number of free tokens must always be greater than or equal to 2 to ensure deadlock free operation.
19:18	Reserved.
17:16	IsocRspTok: isochronous response tokens. Read-write. Cold reset: 0. BIOS: 1.
15:14	IsocPreqTok: isochronous posted request tokens. Read-write. Cold reset: 0. BIOS: 0.
13:12	IsocReqTok: isochronous request tokens. Read-write. Cold reset: 0. BIOS: 1.
11:10	DnRspTok: downstream response tokens. Read-write. Cold reset: 1. BIOS: 1.
9:8	UpRspTok: upstream response tokens. Read-write. Cold reset: 3. BIOS: 1.
7:6	DnPreqTok: downstream posted request tokens. Read-write. Cold reset: 1. BIOS: 1.
5:4	UpPreqTok: upstream posted request tokens. Read-write. Cold reset: 1. BIOS: 1.
3:2	DnReqTok: downstream request tokens. Read-write. Cold reset: 1. BIOS: 1.
1:0	UpReqTok: upstream request tokens. Read-write. Cold reset: 3. BIOS: 1.

D18F3x144 MCT to XCS Token Count

See [D18F3x140](#).

Bits	Description
31:8	Reserved.
7:4	ProbeTok: probe tokens. Read-write. Cold reset: 7h. BIOS: 7h.
3:0	RspTok: response tokens. Read-write. Cold reset: 7h. BIOS: 7h.

D18F3x148 Link to XCS Token Count

See [D18F3x140](#).

Bits	Description
31:30	FreeTok[3:2]: free tokens. See: FreeTok[1:0].
29	Reserved.
28	IsocRspTok1: isochronous response tokens sublink 1. Read-write. Cold reset: 0. BIOS: 0.
27	Reserved.
26	IsocPreqTok1: isochronous posted request tokens sublink 1. Read-write. Cold reset: 0. BIOS: 0.
25	Reserved.
24	IsocReqTok1: isochronous request tokens sublink 1. Read-write. Cold reset: 0. BIOS: 0.
23:22	ProbeTok1: probe tokens sublink 1. Read-write. Cold reset: 0. BIOS: 0.
21:20	RspTok1: response tokens sublink 1. Read-write. Cold reset: 0. BIOS: 0.
19:18	PReqTok1: posted request tokens sublink 1. Read-write. Cold reset: 0. BIOS: 0.
17:16	ReqTok1: request tokens sublink 1. Read-write. Cold reset: 0. BIOS: 0.
15:14	FreeTok[1:0]: free tokens. Read-write. Cold reset: 0000b. FreeTok[3:0] = {FreeTok[3:2], FreeTok[1:0]}. BIOS: 0.
13:12	IsocRspTok0: isochronous response tokens sublink 0. Read-write. Cold reset: 0. BIOS: 0.
11:10	IsocPreqTok0: isochronous posted request tokens sublink 0. Read-write. Cold reset: 0. BIOS: 1. See D18F0x6C[ApplyIsocModeEnNow] .
9:8	IsocReqTok0: isochronous request tokens sublink 0. Read-write. Cold reset: 0. BIOS: 1.
7:6	ProbeTok0: probe tokens sublink 0. Read-write. Cold reset: 2. BIOS: 0.
5:4	RspTok0: response tokens sublink 0. Read-write. Cold reset: 2. BIOS: 2.
3:2	PReqTok0: posted request tokens sublink 0. Read-write. Cold reset: 2. BIOS: 2.
1:0	ReqTok0: request tokens sublink 0. Read-write. Cold reset: 2. BIOS: 2.

D18F3x17C Extended Freelist Buffer Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use [D18F0x6C\[RlsLnkFullTokCntImm\]](#) for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

Bits	Description
------	-------------

31:4	Reserved.
3:0	SPQPrbFreeCBC: XBAR to SRI Probe command buffer freelist. Cold reset: 0h. BIOS: 4h. Read-write.

D18F3x180 Extended NB MCA Configuration

Reset: 0000_0000h. This register is an extension of [D18F3x44 \[MCA NB Configuration\]](#).

Bits	Description
31:27	Reserved.
26	ChgUcToCeEn: change uncorrectable error to correctable error enable. Read-write. 1=The status of uncorrectable errors is changed to appear as correctable errors; D18F3x4C[UC, PCC] are cleared and a machine check exception will not be raised. For uncorrectable ECC errors, D18F3x4C[UECC] is cleared and D18F3x4C[CECC] is set. This field is intended for debug observability.
25	Reserved.
24	McaLogErrAddrWdtErr: log error address on WDT errors. Read-write. BIOS: 1. 1=When a watchdog timeout error occurs (see D18F3x40[WDTRptEn]), the associated address is logged and D18F3x4C[AddrV] is set. 0=When a watchdog timeout error occurs, NB state information is saved and D18F3x4C[AddrV] is cleared. See D18F3x50 for details on saved information.
23	Reserved.
22	SyncFloodOnTblWalkErr: sync flood on table walk error. Read-write. BIOS: 1. 1=A sync flood is generated when the GART table walker encounters an uncorrectable error. A machine check exception is generated independent of the state of this bit.
21	SyncFloodOnCpuLeakErr: sync flood on CPU leak error. Read-write. BIOS: 1. 1=A sync flood is generated when one of the cores encounters an uncorrectable error which cannot be contained to the process on the core.
20	SyncFloodOnL3LeakErr: sync flood on L3 cache leak error. Read-write. BIOS: 1. 1=A sync flood is generated when the L3 cache encounters an uncorrectable error which cannot be contained to the process on one core.
19	PwP2pDatErrRmtPropDis: posted write for remote peer-to-peer data error propagation disable. Read-write. 1= A peer-to-peer posted write with a data error is not propagated to the target IO link chain if the target IO link chain is not attached to the local node (the same node as the source IO link chain). Instead, the write is dropped by the host bridge. This bit can be used in conjunction with DatWrErrDeferEn to cause a machine check exception and SyncFloodOnDeferErrToIO to cause a sync flood.
18	PwP2pDatErrLclPropDis: posted write for local peer-to-peer data error propagation disable. Read-write. 1=A peer-to-peer posted write with a data error is not propagated to the target IO link chain if the target IO link chain is attached to the local node (the same node as the source IO link chain). Instead, the write is dropped by the host bridge. This bit can be used in conjunction with DatWrErrDeferEn to cause a machine check exception and SyncFloodOnDeferErrToIO to cause a sync flood.
17:11	Reserved.
10	Reserved.
9	SyncFloodOnUCNbAry: sync flood on uncorrectable NB array error. Read-write. BIOS: 1. 1=Enables sync flood on detection of an uncorrectable error in a NB array.

8	SyncFloodOnProtocolErr: sync flood on link protocol error. Read-write. BIOS: 1. 1=Enables sync flood on detection of a protocol error.
7	SyncFloodOnTgtAbortErr. Read-write. BIOS: 1. 1=Enable sync flood on generated or received link responses that indicate target aborts.
6	SyncFloodOnDatErr. Read-write. BIOS: 1. 1=Enable sync flood on generated or received link responses that indicate data error.
5	DisPciCfgCpuMstAbortRsp. Read-write. BIOS: 1. 1=For master abort responses to CPU-initiated configuration accesses, disables MCA error reporting and generation of an error response to the core. It is recommended that this bit be set in order to avoid MCA exceptions being generated from master aborts for PCI configuration accesses, which are common during device enumeration.
4	ChgMstAbortToNoErr. Read-write. 1=Signal no errors instead of master abort in link response packets to IO devices on detection of a master abort condition. When ChgMstAbortToNoErr and D18F3x44[IoMstAbortDis] are both set, ChgMstAbortToNoErr takes precedence.
3	ChgDatErrToTgtAbort. Read-write. 1=Signal target abort instead of data error in link response packets to IO devices (for Gen1 link compatibility).
2	WDTCntSel[3]: watchdog timer count select bit[3]. Read-write. See D18F3x44[WDTCntSel] .
1:0	Reserved.

D18F3x190 Downcore Control

Cold reset: 0000_0000h. See [2.4.3 \[Processor Cores and Downcore\]](#) and [2.4.3.1 \[Software Downcore using D18F3x190\[DisCore\]\]](#).

Bits	Description
31:0	<p>DisCore. Read-write; reset-applied. 1=Disable core. 0=Core enabled. [0]=Core 0; ...; [N]=Core N.</p> <ul style="list-style-type: none"> Once a core has been removed by D18F3x190[DisCore]=1, it cannot be added back without a cold reset. E.g. Software may only set DisCore bits, never clear them. The most significant bit N is the number of cores enabled at cold reset and is indicated by the value of CPUID Fn8000_0008_ECX[NC] at cold reset. The most significant bit N and the core ID significance of DisCore is not affected by the value of DisCore followed by a warm-reset. <ul style="list-style-type: none"> E.g. If core 2 is disabled by DisCore[3:0]=0100b followed by a warm reset, then the new core 2 is the old core 3. If the new core 2 needs to then be disabled then DisCore[3:0]=1100b followed by a warm reset. All bits greater than bit N are reserved.

D18F3x1A0 Core to NB Buffer Count

Out of cold reset, the processor allocates a minimal number of buffers that is smaller than the default values in the register. BIOS must use [D18F0x6C\[RIsLnkFullTokCntImm\]](#) for the values in the register to take effect. This is necessary even if the values are unchanged from the default values.

Table 180: Buffer Count Definitions

Term	Definition
NumOfCompUnits	The number of compute units for which at least 1 core is enabled. NumOfCompUnits = COUNT(D18F5x80[Enabled]).

- The following buffer allocations rules must be satisfied:

- CpuCmdBufCnt >= 2.

Bits	Description
31:18	Reserved.
17:16	CpuToNbFreeBufCnt. Read-write. Cold reset: 01b. BIOS: 11b. Provides the number of tokens which can be released to each compute unit from the freelist pool. This field can be updated at any time by BIOS and does not require a warm reset to take effect.
15	Reserved.
14:12	Reserved. Cold reset: 4h.
11:9	Reserved.
8:4	Reserved. Cold reset: Product-specific.
3	Reserved.
2:0	CpuCmdBufCnt: CPU to SRI command buffer count. Read-write; reset-applied. Each compute unit is allocated the number of buffers specified by this field. Cold reset: 2.

D18F3x1CC IBS Control

Reset: 0000_0000h. [MSRC001_103A](#) is an alias of [D18F3x1CC](#). [D18F3x1CC](#) is programmed by BIOS; The OS reads the LVT offset from [MSRC001_103A](#).

Bits	Description						
31:9	Reserved.						
8	LvtOffsetVal: local vector table offset valid. Read-write. BIOS: 1. 1=The offset in LvtOffset is valid. 0=The offset in LvtOffset is not valid and IBS interrupt generation is disabled.						
7:4	Reserved.						
3:0	LvtOffset: local vector table offset. Read-write. BIOS: 0h. Specifies the address of the IBS LVT entry in the APIC registers. See APIC[530:500] . <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>3h-0h</td> <td>LVT address = <500h + LvtOffset<<4></td> </tr> <tr> <td>Fh-4h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	3h-0h	LVT address = <500h + LvtOffset<<4>	Fh-4h	Reserved
Bits	Description						
3h-0h	LVT address = <500h + LvtOffset<<4>						
Fh-4h	Reserved						

D18F3x1FC Product Information Register 1

Bits	Description
31	Reserved.
30	Fp2PcieGen2Sup. IF (CPUID Fn8000_0001_EBX [PkgType] != 0h) THEN Reserved. ELSE Value: Product-specific. 0=The FP2 processor link only supports Gen1 mode. 1=The FP2 processor link supports Gen2 mode. ENDIF.
29	EnDcqChgPriToHigh. Value: Product-specific.
28:25	DllProcessFreqCtlIndex2Rate50 [3:0]. Value: Product-specific.
24	SWDLLCapTableEn. Value: Product-specific.
23	ForceSmcCheckFlwStDis. Value: Product-specific. See MSRC001_102D [ForceSmcCheckFlwStDis].
22	DiDtCfg5. Value: Product-specific. See MSRC001_1028 [DiDtCfg5].

21	Reserved.
20	EnCstateBoostBlockCC6Exit . Value: 0. See D18F5x88 [EnCstateBoostBlockCC6Exit].
19:17	DiDtCfg4 . Value: Product-specific. See MSRC001_1028 [DiDtCfg4].
16	Reserved.
15:14	DiDtCfg2 . Value: Product-specific. See MSRC001_1028 [DiDtCfg2].
13:6	DiDtCfg1 . Value: Product-specific. See MSRC001_1028 [DiDtCfg1].
5:1	DiDtCfg0 . Value: Product-specific. See MSRC001_1028 [DiDtCfg0].
0	DiDtMode . Value: Product-specific. See MSRC001_1028 [DiDtMode].

3.12 Device 18h Function 4 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space].

D18F4x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1404h.
15:0	VendorID: vendor ID. Read-only. Value: 1022h.

D18F4x04 Status/Command

Bits	Description
31:16	Status. Read-only. Reset: 0000_0000_000X_0000b. Only Status[4] may be set to indicate the existence of a PCI-defined capability block. 0=No supported links are ungangled. 1=At least one link may be ungangled, in which case there is a capability block associated with sublink one of the link in this function.
15:0	Command. Read-only. Value: 0000h.

D18F4x08 Class Code/Revision ID

Reset: 0600_0000h.

Bits	Description
31:8	ClassCode. Read-only. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only.

D18F4x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:0	HeaderTypeReg. Read-only. These bits are fixed at their default values. The header type field indicates that there are multiple functions present in this device.

D18F4x34 Capabilities Pointer

Bits	Description
31:8	Reserved.
7:0	CapPtr: capabilities pointer. Read-only. Value: 00h.

D18F4x108 TDP Limit 1

Bits	Description
31:0	Reserved.

D18F4x10C TDP Limit 2

Updated-by-hardware.

Bits	Description
31:12	Reserved.
11:0	NodeTdpLimit. Read-write. Reset: Product-specific. Specifies the maximum allowed sum of TDPs from all cores on a node. If the consumed power exceeds the NodeTdpLimit, a P-state limit is applied to all cores on the processor to reduce the power consumption so that it remains within the TDP limit. If D18F4x15C[BoostLock]=1 , NodeTdpLimit can only be written with values that are less than or equal to the reset value. Attempts to write an invalid value are ignored. See .

D18F4x110 Sample and Residency Timers

Bits	Description
31:21	Reserved.
20:13	MinResTmr: minimum residency timer. IF D18F4x15C[BoostLock] THEN Read-only. ELSE Read-write. ENDIF. Cold reset: Product-specific. Specifies the minimum amount of time required between TDP-initiated P-state transitions. The minimum amount of time is defined as $\text{MinResTmr} * \text{CSampleTimer} * 5.12\mu\text{s}$.
12	Reserved.
11:0	CSampleTimer. Read-write. Cold reset: 0. BIOS: 002h. Specifies the value that the internal CSampleTimer counter must increment to before expiring. When the internal CSampleTimer counter expires, it is reset to 0. See 2.5.3.1.1 [Application Power Management (APM)] .

D18F4x11[C:8] C-state Control

[D18F4x11\[C:8\]](#) consist of three identical 16-bit registers, one for each C-state Action Field (CAF) associated with an IO address that is read to enter C-states. Refer to [2.5.3.2 \[Core C-states\]](#).

- [D18F4x118\[15:0\]](#) specifies the actions attempted by the core when software reads from the IO address specified by [MSRC001_0073\[CstateAddr\]](#).
- [D18F4x118\[31:16\]](#) specifies the actions attempted by the core when software reads from the IO address specified by [MSRC001_0073\[CstateAddr+1\]](#).
- [D18F4x11C\[15:0\]](#) specifies the actions attempted by the core when software reads from the IO address specified by [MSRC001_0073\[CstateAddr+2\]](#).

D18F4x118 C-state Control 1

Reset: 0000_0000h. Read-write, updated-by-hardware.

Bits	Description
31:30	Reserved.
29	SelfRefrEarly1 . See: SelfRefrEarly0. BIOS: 0.
28	SelfRefr1 . See: SelfRefr0. BIOS: 1.
27	NbClkGate1 . See: NbClkGate0. BIOS: 1.
26	NbPwrGate1 . See: NbPwrGate0. BIOS: 1.
25	PwrOffEnCstAct1 . See: PwrOffEnCstAct0. BIOS: 1.
24	PwrGateEnCstAct1 . See: PwrGateEnCstAct0. BIOS: 1.
23:21	ClkDivisorCstAct1 . See: ClkDivisorCstAct0. BIOS: 000b.
20	Reserved.
19:18	CacheFlushTmrSelCstAct1 . See: CacheFlushTmrSelCstAct0. BIOS: 01b.
17	CacheFlushEnCstAct1 . See: CacheFlushEnCstAct0. BIOS: 1.
16	CpuPrbEnCstAct1 . See: CpuPrbEnCstAct0. BIOS: 1.
15:14	Reserved.
13	SelfRefrEarly0: allow early self-refresh . BIOS: 0. 1=Allow self-refresh while cores in CC1 are waiting for the cache flush timer to expire. 0=Wait for cache flush timer to expire before allowing self-refresh. See 2.5.7.2 [DRAM Self-Refresh] and 2.5.3.2.3.1 [C-state Probes and Cache Flushing] .
12	SelfRefr0: self-refresh . BIOS: 1. 1=Allow DRAM self-refresh while in NB C-states. 0=Prevent DRAM self-refresh while in NB C-states. NbClkGate0 must be programmed to 1 if this field is programmed to 1. See 2.5.7.2 [DRAM Self-Refresh] and 2.5.4.2 [NB C-states] .
11	NbClkGate0: NB clock-gating . BIOS: 1. 1=Allow clock-gating of the NB. 0=Prevent clock-gating of the NB. See 2.5.4.2 [NB C-states] .
10	NbPwrGate0: NB power-gating . BIOS: 1. 1=Allow power-gating of the NB. 0=Prevent power-gating of the NB. See 2.5.4.2 [NB C-states] .
9	PwrOffEnCstAct0: power off enable . BIOS: 1. 1=Package power off enable. CacheFlushEnCstAct0 is required to be set if this bit is set. PwrGateEnCstAct0 is required to be set if this bit is set. See 2.5.3.2.3.4 [Package C6 (PC6) State] .
8	PwrGateEnCstAct0: power gate enable . BIOS: 1. 1=Core power gating is enabled. CacheFlushEnCstAct0 is required to be set if this bit is set. See 2.5.3.2.3.3 [Core C6 (CC6) State] .

7:5	<p>ClkDivisorCstAct0: clock divisor. BIOS: 000b. Specifies the core clock frequency while in the low-power state before the caches are flushed. This divisor is relative to the current FID frequency, or:</p> <ul style="list-style-type: none"> • 100 MHz * (10h + MSRC001_00[6B:64][CpuFid]) of the current P-state specified by MSRC001_0063[CurPstate]. <p>If MSRC001_00[6B:64][CpuDid] of the current P-state indicates a divisor that is deeper than specified by this field, then no frequency change is made when entering the low-power state associated with this register.</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;"><u>Bits</u></th> <th style="text-align: left;"><u>Description</u></th> <th style="text-align: left;"><u>Bits</u></th> <th style="text-align: left;"><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>/1</td> <td>100b</td> <td>/16</td> </tr> <tr> <td>001b</td> <td>/2</td> <td>101b</td> <td>/128</td> </tr> <tr> <td>010b</td> <td>/4</td> <td>110b</td> <td>/512</td> </tr> <tr> <td>011b</td> <td>/8</td> <td>111b</td> <td>Turn off clocks.</td> </tr> </tbody> </table> <p>See CacheFlushTmrSelCstAct0.</p>	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>	000b	/1	100b	/16	001b	/2	101b	/128	010b	/4	110b	/512	011b	/8	111b	Turn off clocks.
<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>																		
000b	/1	100b	/16																		
001b	/2	101b	/128																		
010b	/4	110b	/512																		
011b	/8	111b	Turn off clocks.																		
4	Reserved.																				
3:2	<p>CacheFlushTmrSelCstAct0: cache flush timer select. BIOS: 10b. Specifies the timer to use for cache flush if (D18F4x128[CoreCstateMode]==0).</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;"><u>Bits</u></th> <th style="text-align: left;"><u>Cache flush timer</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>0 us</td> </tr> <tr> <td>01b</td> <td>D18F3xDC[CacheFlushOnHaltTmr]</td> </tr> <tr> <td>10b</td> <td>D18F4x128[CacheFlushTmr]</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table> <p>Each compute unit has one timer that is shared by both cores. D18F3xDC[CacheFlushOnHaltCtl] specifies the core clock divisor to use after the caches are flushed. Writing values greater than 10b result in 10b. See CacheFlushEnCstAct0 and CpuPrbEnCstAct0.</p>	<u>Bits</u>	<u>Cache flush timer</u>	00b	0 us	01b	D18F3xDC[CacheFlushOnHaltTmr]	10b	D18F4x128[CacheFlushTmr]	11b	Reserved										
<u>Bits</u>	<u>Cache flush timer</u>																				
00b	0 us																				
01b	D18F3xDC[CacheFlushOnHaltTmr]																				
10b	D18F4x128[CacheFlushTmr]																				
11b	Reserved																				
1	<p>CacheFlushEnCstAct0: cache flush enable. BIOS: 1. 1=Cache flush enable if (D18F4x128[CoreCstateMode]==0). The cache flush timer starts counting when the C-state is entered. See CacheFlushTmrSelCstAct0 and 2.5.3.2.3.1 [C-state Probes and Cache Flushing].</p>																				
0	<p>CpuPrbEnCstAct0: core direct probe enable. BIOS: 1. Specifies how probes are handled while in the low-power state. 0=When the probe request comes into the NB, the core clock is brought up to the COF (based on the current P-state), all outstanding probes are completed, the core waits for a hysteresis time based on D18F3xD4[ClkRampHystSel], and then the core clock is brought down to the frequency specified by ClkDivisorCstAct0. 1=The core clock does not change frequency; the probe is handled at the frequency specified by ClkDivisorCstAct0; this may only be set if:</p> <ul style="list-style-type: none"> • ClkDivisorCstAct0 specifies a divide-by 1, 2, 4, 8, or 16 and NbCof <= 3.2 GHz • ClkDivisorCstAct0 specifies a divide-by 1, 2, 4, or 8 and NbCof >= 3.4 GHz <p>This bit also specifies functionality of the timer used for cache flushing. See CacheFlushTmrSelCstAct0.</p> <ul style="list-style-type: none"> • If CpuPrbEnCstAct0=0 and D18F3xDC[IgnCpuPrbEn]=0, only the time when the core is in a non-C0 state and has its clocks ramped up to service probes is counted. • If CpuPrbEnCstAct0=1 or D18F3xDC[IgnCpuPrbEn]=1, all of the time the core is in a non-C0 state is counted. 																				

D18F4x11C C-state Control 2

Reset: 0000_0000h. Read-write.

Bits	Description
31:14	Reserved.
13	SelfRefrEarly2. See: D18F4x11[C:8][SelfRefrEarly0] .
12	SelfRefr2. See: D18F4x11[C:8][SelfRefr0] .
11	NbClkGate2. See: D18F4x11[C:8][NbClkGate0] .
10	NbPwrGate2. See: D18F4x11[C:8][NbPwrGate0] .
9	PwrOffEnCstAct2. See: D18F4x11[C:8][PwrOffEnCstAct0] .
8	PwrGateEnCstAct2. See: D18F4x11[C:8][PwrGateEnCstAct0] .
7:5	ClkDivisorCstAct2. See: D18F4x11[C:8][ClkDivisorCstAct0] .
4	Reserved.
3:2	CacheFlushTmrSelCstAct2. See: D18F4x11[C:8][CacheFlushTmrSelCstAct0] .
1	CacheFlushEnCstAct2. See: D18F4x11[C:8][CacheFlushEnCstAct0] .
0	CpuPrbEnCstAct2. See: D18F4x11[C:8][CpuPrbEnCstAct0] .

D18F4x124 C-state Interrupt Control

Reset: 0000_0000h. Read-write. See [2.5.3.2.4.2 \[Timer Tick Monitor\]](#).

Bits	Description										
31:27	Reserved.										
26:23	IntMonPC6Limit: interrupt monitor package C6 limit. Read-write. Specifies the threshold for disallowing access to the PC6 state. <table border="1"> <thead> <tr> <th>Bits</th> <th>Threshold</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>PC6 entry disallowed</td> </tr> <tr> <td>1001b-0001b</td> <td>IntMonPC6Limit * 1ms</td> </tr> <tr> <td>1110b-1010b</td> <td>((5 * IntMonPC6Limit) - 40) * 1ms</td> </tr> <tr> <td>1111b</td> <td>40ms</td> </tr> </tbody> </table>	Bits	Threshold	0000b	PC6 entry disallowed	1001b-0001b	IntMonPC6Limit * 1ms	1110b-1010b	((5 * IntMonPC6Limit) - 40) * 1ms	1111b	40ms
Bits	Threshold										
0000b	PC6 entry disallowed										
1001b-0001b	IntMonPC6Limit * 1ms										
1110b-1010b	((5 * IntMonPC6Limit) - 40) * 1ms										
1111b	40ms										
22	IntMonPC6En: interrupt monitor package C6 enable. Read-write. Specifies whether the interrupt monitor is enabled for the PC6 state. 0=Disabled. 1=Enabled.										
21:0	Reserved.										

D18F4x128 C-state Policy Control 1

Reset: 0000_0000h.

Bits	Description
31	CstateMsgDis: C-state messaging disable. Read-write. BIOS: 1. Specifies whether any messages are sent to the FCH when a core enters or exits a C-state. 0=Messages are sent. 1=Messages are not sent. See 2.5.3.2.4.1 [FCH Messaging] .
30:21	Reserved.
20:18	CacheFlushSucMonThreshold: cache flush success monitor threshold. Read-write. BIOS: 111b. Flush the caches immediately if cache flushing is enabled and the cache flush success monitor count > CacheFlushSucMonThreshold. A value of 0 disables the cache flush success monitor. See D18F4x11[C:8][CacheFlushEn] .

17:12	Reserved.						
11:5	<p>CacheFlushTmr: cache flush timer. Read-write. BIOS: 14h. Specifies how long each core needs to stay in a C-state before it flushes its caches. See D18F4x11[C:8][CacheFlushTmrSel].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>$\leq 5.12 \text{ us}$</td> </tr> <tr> <td>7Fh-01h</td> <td>$(\langle \text{CacheFlushTmr} \rangle * 10.24\text{us}) - 5.12\text{us} \leq \text{Time} \leq \langle \text{CacheFlushTmr} \rangle * 10.24 \text{ us}$</td> </tr> </tbody> </table>	Bits	Description	00h	$\leq 5.12 \text{ us}$	7Fh-01h	$(\langle \text{CacheFlushTmr} \rangle * 10.24\text{us}) - 5.12\text{us} \leq \text{Time} \leq \langle \text{CacheFlushTmr} \rangle * 10.24 \text{ us}$
Bits	Description						
00h	$\leq 5.12 \text{ us}$						
7Fh-01h	$(\langle \text{CacheFlushTmr} \rangle * 10.24\text{us}) - 5.12\text{us} \leq \text{Time} \leq \langle \text{CacheFlushTmr} \rangle * 10.24 \text{ us}$						
4:2	HaltCstateIndex. Read-write. BIOS: 0. Specifies the IO-based C-state that is invoked by a HLT instruction.						
1	CoreCstatePolicy. Read-write. BIOS: 0. Specifies how the processor arbitrates voltage and frequency when different non-C0 C-state requests are received on each core in a compute unit. 0=Transition both cores to the shallower C-state request. 1=Transition both cores to the deeper C-state request. For instance, if core 0 gets a request to go to C2 and core 1 gets a request to go to C1, hardware looks at the setting of CoreCstatePolicy. If CoreCstatePolicy is programmed to 0, the processor sends both cores to C1. If CoreCstatePolicy is programmed to 1, the processor sends both cores to C2. BIOS should program this field to the same value in all nodes of a multi-node processor. See also 2.5.2.2 [Dependencies Between Subcomponents on VDDNB] .						
0	Reserved.						

D18F4x13C SMU P-state Control

Updated-by-hardware.

Bits	Description
31:4	Reserved.
3:1	SmuPstateLimit. Read-only; updated-by-hardware. Reset: 0.
0	SmuPstateLimitEn. Read-only; updated-by-hardware. Reset: 0.

D18F4x15C Core Performance Boost Control

Updated-by-hardware.

Bits	Description
31	<p>BoostLock. Read-only. Reset: Product-specific. Specifies whether the following registers are read-write, read-only, or have special requirements related to writability. See individual register definitions for details.</p> <ul style="list-style-type: none"> • MSRC001_00[6B:64][CpuFid, CpuDid, CpuVid]. • D18F4x10C[NodeTdpLimit]. • D18F4x110[MinResTmr]. • D18F4x15C[NumBoostStates]. • D18F4x16C[CstateCnt, CstateBoost].
30:8	Reserved.
7	ApmMasterEn: APM master enable. Read-write. Reset: 0. BIOS: IF(D18F4x15C[NumBoostStates] ==0) THEN 0. ELSE 1. ENDIF. 1=Enables the ability to turn on features associated with APM when used in conjunction with the individual feature enable bits. See 2.5.3.1.1 [Application Power Management (APM)] .
6:5	Reserved.

4:2	NumBoostStates: number of boosted states. IF (D18F4x15C[BoostLock] ApmMasterEn) THEN Read-only. ELSE Read-write. ENDIF. Reset: Product-specific. Specifies the number of P-states that are considered boosted P-states. See 2.5.3.1.1 [Application Power Management (APM)].										
1:0	BoostSrc: boost source. Read-write. Reset: 0. BIOS: 2.5.3.1.7. Specifies whether CPB is enabled or disabled. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Boosting disabled</td> </tr> <tr> <td>01b</td> <td>Boosting enabled</td> </tr> <tr> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	Boosting disabled	01b	Boosting enabled	10b	Reserved	11b	Reserved
Bits	Description										
00b	Boosting disabled										
01b	Boosting enabled										
10b	Reserved										
11b	Reserved										

D18F4x164 Fixed Errata

Bits	Description
31:0	FixedErrata. Value: Product-specific.. See the <i>Revision Guide for AMD Family 15h Models 10h-1Fh Processors</i> for the definition of this field.

D18F4x16C APM TDP Control

Bits	Description
31:12	Reserved.
11:9	CstateCnt: C-state count. IF D18F4x15C[BoostLock] THEN Read-only. ELSE Read-write. ENDIF. Reset: Product-specific. Specifies the number of compute units that must be in CC6 before a transition can occur to a boosted P-state that is higher performance than the P-state specified by CstateBoost. A value of 0 disables access to P-states above CstateBoost.
8:6	CstateBoost. Read-write. Reset: Product-specific. Specifies the P-state which requires the number of compute units specified in CstateCnt to be in CC6 before a transition to a higher performance (lower numbered) boosted P-state is allowed. CstateBoost must be less than or equal to D18F4x15C[NumBoostStates] otherwise undefined behavior results. If D18F4x15C[BoostLock]=1, CstateBoost can only be written with values that are greater than or equal to the reset value. Attempts to write values less than the reset value are ignored. A value of 0 indicates that the C-state boost feature is not supported. This field uses hardware P-state numbering. See 2.5.3.1.2.2 [Hardware P-state Numbering].
5	Reserved.
4	ApmTdpLimitIntEn: APM TDP limit interrupt enable. Read-write. Reset: 0. 1=Enables the generation of an interrupt using APIC330 of each core when TDP changes.
3:0	Reserved.

3.13 Device 18h Function 5 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.7 [Configuration Space].

D18F5x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1405h.
15:0	VendorID: vendor ID. Read-only. Value: 1022h.

D18F5x04 Status/Command

Bits	Description
31:16	Status. Read-only. Value: 0000h.
15:0	Command. Read-only. Value: 0000h.

D18F5x08 Class Code/Revision ID

Bits	Description
31:8	ClassCode. Read-only. Value: 060000h. Provides the host bridge class code as defined in the PCI specification.
7:0	RevID: revision ID. Read-only. Value: 00h.

D18F5x0C Header Type

Reset: 0080_0000h.

Bits	Description
31:0	HeaderTypeReg. Read-only. These bits are fixed at their default values. The header type field indicates that there are not multiple functions present in this device.

D18F5x34 Capabilities Pointer

Bits	Description
31:8	Reserved.
7:0	CapPtr: capabilities pointer. Read-only. Value: 00h.

D18F5x[70,60,50,40] Northbridge Performance Event Select Low

Bits	Description
31:0	MSRC001_024[6,4,2,0][31:0] is an alias of D18F5x[70,60,50,40] .

D18F5x[74,64,54,44] Northbridge Performance Event Select High

Bits	Description
31:0	MSRC001_024[6,4,2,0][63:32] is an alias of D18F5x[74,64,54,44].

D18F5x[78,68,58,48] Northbridge Performance Event Counter Low

Bits	Description
31:0	MSRC001_024[7,5,3,1][31:0] is an alias of D18F5x[78,68,58,48].

D18F5x[7C,6C,5C,4C] Northbridge Performance Event Counter High

Bits	Description
31:0	MSRC001_024[7,5,3,1][63:32] is an alias of D18F5x[7C,6C,5C,4C].

D18F5x80 Compute Unit Status

Read-only. See 2.4.3 [Processor Cores and Downcoring]. Reset: Product-specific.

Software associates core ID to the cores of the compute units according to the following table. All combinations not listed are reserved.

Table 181: D18F5x80[Enabled, DualCore] Definition

Enabled	DualCore	Definition
1h	1h	1 compute unit is enabled; both cores of the compute unit are enabled.
3h	3h	2 compute units are enabled; both cores of each compute unit are enabled.

Bits	Description										
31:20	Reserved.										
19:16	DualCore: both cores of a compute unit are enabled. 1=Both cores of a compute unit are enabled. See Table 181 [D18F5x80[Enabled, DualCore] Definition]. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Compute unit 0</td> </tr> <tr> <td>[1]</td> <td>Compute unit 1</td> </tr> <tr> <td>[2]</td> <td>Reserved</td> </tr> <tr> <td>[3]</td> <td>Reserved</td> </tr> </tbody> </table>	Bit	Description	[0]	Compute unit 0	[1]	Compute unit 1	[2]	Reserved	[3]	Reserved
Bit	Description										
[0]	Compute unit 0										
[1]	Compute unit 1										
[2]	Reserved										
[3]	Reserved										
15:4	Reserved.										
3:0	Enabled: at least one core of a compute unit is enabled. 1=At least one core is enabled in a compute unit. See Table 181 [D18F5x80[Enabled, DualCore] Definition]. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Compute unit 0</td> </tr> <tr> <td>[1]</td> <td>Compute unit 1</td> </tr> <tr> <td>[2]</td> <td>Reserved</td> </tr> <tr> <td>[3]</td> <td>Reserved</td> </tr> </tbody> </table>	Bit	Description	[0]	Compute unit 0	[1]	Compute unit 1	[2]	Reserved	[3]	Reserved
Bit	Description										
[0]	Compute unit 0										
[1]	Compute unit 1										
[2]	Reserved										
[3]	Reserved										

D18F5x84 Northbridge Capabilities 2

Read-only. Value: Product-specific. Unless otherwise specified, 1=The feature is supported by the processor; 0=The feature is not supported.

Bits	Description
31:29	Reserved.
28:24	DdrMaxRateEnf: enforced maximum DDR rate. See: DdrMaxRate. Specifies the maximum DRAM data rate that the processor is designed to support. Writes to D18F2x94_dct[1:0][MemClkFreq] that specify a frequency greater than specified by DdrMaxRateEnf will result in the D18F2x94_dct[1:0][MemClkFreq] being set to DdrMaxRateEnf.
23:21	Reserved.
20:16	DdrMaxRate: maximum DDR rate. Specifies the maximum DRAM data rate that the processor is designed to support. See: Table 132 [Memory Clock Frequency Value Definition] ; except 00h is defined as no limit. See D18F2x94_dct[1:0][MemClkFreq] and D18F5x84[DdrMaxRateEnf] .
15:14	Reserved.
13:12	DctEn: DCT enabled. Specifies which DCT controllers are enabled. [0]=DCT0. [1]=DCT1. 1=Enabled. 0=Disabled.
11:8	Reserved.
7:0	CmpCap: CMP capable. Number of enabled cores on the node is CmpCap+1. CmpCap does not reflect cores disabled by D18F3x190[DisCore] .

D18F5x88 Northbridge Configuration 4 (NB_CFG4)

Bits	Description
31:25	Reserved.
24	DisHbNpReqBusLock. Read-write. Reset: 0. BIOS: 1. 0=While bus locks are in progress, all non-posted commands from I/O, including atomics, are blocked until the core has completed the locked transaction and releases the bus. 1=All non-posted commands except atomics do not honor bus locks and are allowed to proceed. This bit may be set to achieve better DMA performance in the presence of bus locks.
23:20	Reserved.
19	EnCpuInSWP0DctHint. Read-write. Reset: 0. The processor allocates additional memory bandwidth to the GPU when all cores are not in P0. This bit controls whether software P0 or hardware P0 is used. 1=Use software P0. 0=Use hardware P0. See 2.5.3.1.2 [Core P-state Naming and Numbering] .
18	EnCstateBoostBlockCC6Exit. Read-write. Reset: 0. BIOS: D18F3x1FC[EnCstateBoostBlockCC6Exit] . If (EnCstateBoostBlockCC6Exit==1), cores cannot exit CC6 until VDD is less than or equal to the voltage of the P-state indexed by D18F4x16C[CstateBoost] .
17:3	Reserved.

2	IntStpClkHaltExitEn. Read-write. Reset: 0. BIOS: 1. 1=The Northbridge sends out HALT_EXIT messages received from the core due to internal StpClk interrupts. 0=The Northbridge ignores HALT_EXIT messages received from the core due to internal StpClk interrupts. This bit should always be set.
1:0	Reserved.

D18F5xE0 Processor TDP Running Average

Bits	Description
31:4	Reserved.
3:0	RunAvgRange: running average range. Read-write. Reset: 0. BIOS: 2h. Specifies the interval over which the processor averages power consumption estimates from the cores for boosting. Time interval = $2^{(\text{RunAvgRange} + 1)} * \text{FreeRunSampleTimer}$ rate. A value of 0 disables the TDP running average accumulator capture function. See 2.5.3.1.1 [Application Power Management (APM)].

D18F5x128 Clock Power/Timing Control 3

Bits	Description												
31:23	Reserved.												
22	NbPllPwrDwnRegEn: NB PLL power down . Read-write. Cold reset: Product-specific. 1=The NB PLL is powered down when the NB is power gated and DRAM is placed into self-refresh (see 2.5.4.2 [NB C-states]). 0=The NB PLL is not powered down during NB C-states.												
21	PC6Vid[7]. Read-write. Cold reset: Product-specific. See PC6Vid[6:0].												
20:16	Reserved.												
15	CC6PwrDwnRegEn: CC6 power down regulator enable. Read-write. Cold reset: Product-specific. 1=Power down the VDDA regulator on CC6 entry. See PllRegTime.												
14	PC6PwrDwnRegEn: PC6 power down regulator enable. Read-write. Cold reset: Product-specific. 1=Power down the VDDA regulator on PC6 entry. See PllRegTime.												
13:12	PwrGateTmr: power gate timer. Read-write. Cold reset: 01b. BIOS: 01b. Specifies the minimum delay time required from the power gating or ungating of one compute unit to the power gating or ungating of the same compute unit or another compute unit. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>500 ns</td> <td>10b</td> <td>5 us</td> </tr> <tr> <td>01b</td> <td>1 us</td> <td>11b</td> <td>10 us</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	00b	500 ns	10b	5 us	01b	1 us	11b	10 us
Bits	Description	Bits	Description										
00b	500 ns	10b	5 us										
01b	1 us	11b	10 us										
11:10	PllVddOutUpTime. Read-write. Cold Reset: 0. BIOS: 11b. The VDD regulator may be powered down when the processor transitions to PC6. If the regulator is powered down, this field specifies the time required to initialize the core PLL logic once the regulator is powered back up. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>100 ns</td> <td>10b</td> <td>400 ns</td> </tr> <tr> <td>01b</td> <td>200 ns</td> <td>11b</td> <td>800 ns</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	00b	100 ns	10b	400 ns	01b	200 ns	11b	800 ns
Bits	Description	Bits	Description										
00b	100 ns	10b	400 ns										
01b	200 ns	11b	800 ns										
9	FastSlamTimeDown. Read-write. Cold Reset: 0. BIOS: 1. Specifies the time the processor waits for downward voltage transitions to complete. This field only effects transitions from D18F4x16C[CstateBoost] or lower performance P-states. 0=D18F3xD8[VSRampSlamTime] . 1=10 us.												

8:7	<p>PllRegTime: Pll regulator time. Read-write. Cold Reset: 10b. The VDDAregulator may be powered down when the processor transitions to PC6 or CC6. See PC6PwrDwnRegEn and CC6PwrDwnRegEn. If CC6PwrDwnRegEn=1, the VDDA regulator is powered down during CC6. If PC6PwrDwnRegEn=1, the VDDA regulator is powered down during PC6. If the VDDA regulator is powered down during CC6 and the core transitions from CC6 to PC6, the regulator remains powered down during PC6 regardless of the PC6PwrDwnRegEn setting. This field specifies the time required for the VDDA regulator to power back up and initialize the core PLL logic that is powered by the VDDA regulator.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Reserved.</td> <td>10b</td> <td>1.5 us</td> </tr> <tr> <td>01b</td> <td>Reserved.</td> <td>11b</td> <td>2.0 us</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>	00b	Reserved.	10b	1.5 us	01b	Reserved.	11b	2.0 us
<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>										
00b	Reserved.	10b	1.5 us										
01b	Reserved.	11b	2.0 us										
6:0	<p>PC6Vid[6:0]: package C6 vid. Read-write. Cold reset: Value varies by product. PC6Vid[7:0] = {PC6Vid[7], PC6Vid[6:0]}. PC6Vid[7:0] specifies the VID driven in the PC6 state. See 2.5.3.2.3.4 [Package C6 (PC6) State] and 2.5.1.3.2 [Low Power Voltages]. Required to be programmed within the range specified by MSRC001_0071[MaxVid and MinVid].</p>												

D18F5x12C Clock Power/Timing Control 4

Cold reset: 0000_000Eh. See the *AMD Serial VID Interface 2.0 (SVI2) Specification*.

Bits	Description																				
31	<p>Svi2CmdBusy. Read-only, updated-by-hardware. 1=SVI2 command in progress. This bit is set by hardware when any SVI2 command is sent to the voltage regulator. Software must wait for this bit to clear to 0 before writing any of the below fields:</p> <ul style="list-style-type: none"> • D18F5x12C[CorePsi1En, CoreLoadLineTrim, CoreOffsetTrim] • D18F5x188[NbPsi1, NbLoadLineTrim, NbOffsetTrim] <p>This bit is cleared by hardware when the SVI2 command is complete. On a voltage change, this bit is cleared when the voltage transition is completed. See 2.5.1.4.1 [Hardware-Initiated Voltage Transitions]. On a PSIx_L change, this bit is cleared as soon as the SVI2 command is sent to the voltage regulator. See 2.5.1.1.1 [SVI2 Features] and 2.5.1.3.1 [PSIx_L Bit].</p>																				
30	<p>WaitVidCompDis: wait VID completion disable. Read-write. 0=Hardware waits for the VOTF complete indicator from the voltage regulator before clearing Svi2CmdBusy or making additional voltage change requests. 1=Hardware clears Svi2CmdBusy 500us after changes to CoreLoadLineTrim, CoreOffsetTrim, or D18F5x188[NbLoadLineTrim, NbOffsetTrim] are made. Hardware clears Svi2CmdBusy and additional voltage changes are allowed after the time specified by 2.5.1.4 [Voltage Transitions].</p>																				
29:6	Reserved.																				
5	<p>CorePsi1En: Core PSI1_L enable. Read-write. BIOS: 1. 0=PSI1_L for VDD is deasserted. 1=PSI1_L for VDD is asserted when all cores are in CC6. See 2.5.3.2.3.4 [Package C6 (PC6) State], 2.5.1.3.1 [PSIx_L Bit], and Svi2CmdBusy.</p>																				
4:2	<p>CoreLoadLineTrim: Core load line trim. Read-write. BIOS: D0F0xBC_xE0104184[SviLoadLineTrimVdd]. CoreLoadLineTrim and NbLoadLineTrim specify a percentage change relative to the initial load line slope for VDD and VDDNB, respectively.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Load line disabled</td> <td>100b</td> <td>+20%</td> </tr> <tr> <td>001b</td> <td>-40%</td> <td>101b</td> <td>+40%</td> </tr> <tr> <td>010b</td> <td>-20%</td> <td>110b</td> <td>+60%</td> </tr> <tr> <td>011b</td> <td>0%</td> <td>111b</td> <td>+80%</td> </tr> </tbody> </table> <p>See Svi2CmdBusy.</p>	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>	000b	Load line disabled	100b	+20%	001b	-40%	101b	+40%	010b	-20%	110b	+60%	011b	0%	111b	+80%
<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>																		
000b	Load line disabled	100b	+20%																		
001b	-40%	101b	+40%																		
010b	-20%	110b	+60%																		
011b	0%	111b	+80%																		

1:0	CoreOffsetTrim: Core offset trim. Read-write. BIOS: D0F0xBC_xE0104184[SviLoadLineOffset-Vdd]. CoreOffsetTrim and NbOffsetTrim specify a voltage offset relative to the initial load line offset for VDD and VDDNB, respectively.			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	00b	Load line offset disabled	10b	0mV
	01b	-25mV	11b	+25mV
	See Svi2CmdBusy.			

D18F5x1[6C:60] Northbridge P-state [3:0]

Cold reset: Product-specific. Each of these registers specify the frequency and voltage associated with each of the NB P-states.

Table 182: Register Mapping for D18F5x1[6C:60]

Register	Function
D18F5x160	NB P-state 0
D18F5x164	NB P-state 1
D18F5x168	NB P-state 2
D18F5x16C	NB P-state 3

The NbVid field is allowed to be different between processors in a multi-processor system. All other fields are required to be programmed to the same value for all processors in the coherent fabric. See 2.5.4.1 [NB P-states] for more information about these registers.

Table 183: NB P-state Definitions

Term	Definition
NBCOF	NB current operating frequency in MHz. $NBCOF = 100 * (D18F5x1[6C:60][NbFid] + 4h) / (2^{D18F5x1[6C:60][NbDid]})$.
NBCOF[0]	NB current operating frequency in MHz for NB P-state 0. $NBCOF[0] = (100 * (D18F5x160[NbFid] + 4h) / (2^{D18F5x160[NbDid]}))$.
NBCOF[1]	NB current operating frequency in MHz for NB P-state 1. $NBCOF[1] = (100 * (D18F5x164[NbFid] + 4h) / (2^{D18F5x164[NbDid]}))$.
NBCOF[2]	NB current operating frequency in MHz for NB P-state 2. $NBCOF[2] = (100 * (D18F5x168[NbFid] + 4h) / (2^{D18F5x168[NbDid]}))$.
NBCOF[3]	NB current operating frequency in MHz for NB P-state 3. $NBCOF[3] = (100 * (D18F5x16C[NbFid] + 4h) / (2^{D18F5x16C[NbDid]}))$.

Bits	Description
31:24	NbIddValue: northbridge current value. Read-write. See NbIddDiv.

23:22	NbIddDiv: northbridge current divisor. Read-write. After reset, NbIddDiv and NbIddValue combine to specify the expected maximum current drawn on the VDDNB power plane at a given VDDNB voltage. These values are intended to be used by 2.5.1.3.1.1 [BIOS Requirements for PSIO_L] . These values are not intended to convey final product power levels. These fields may be subsequently altered by software; they do not affect the hardware behavior. <table border="1"> <thead> <tr> <th><u>NbIddDiv</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>IddValue / 1 A, Range: 0 to 255 A.</td> </tr> <tr> <td>01b</td> <td>IddValue / 10 A, Range: 0 to 25.5 A.</td> </tr> <tr> <td>10b</td> <td>IddValue / 100 A, Range: 0 to 2.55 A.</td> </tr> <tr> <td>11b</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>NbIddDiv</u>	<u>Description</u>	00b	IddValue / 1 A, Range: 0 to 255 A.	01b	IddValue / 10 A, Range: 0 to 25.5 A.	10b	IddValue / 100 A, Range: 0 to 2.55 A.	11b	Reserved.
<u>NbIddDiv</u>	<u>Description</u>										
00b	IddValue / 1 A, Range: 0 to 255 A.										
01b	IddValue / 10 A, Range: 0 to 25.5 A.										
10b	IddValue / 100 A, Range: 0 to 2.55 A.										
11b	Reserved.										
21	NbVid[7]. Read-write. See NbVid[6:0].										
20:19	Reserved.										
18	MemPstate: Memory P-state. Read-write. 1=The Northbridge P-state specified by this register maps to memory P-state 1. 0=The Northbridge P-state specified by this register maps to memory P-state 0. Memory P-states may be globally disabled by programming D18F5x170[MemPstateDis] . See 2.5.7.1 [Memory P-states] .										
17	Reserved.										
16:10	NbVid[6:0]: Northbridge VID. Read-write. NbVid[7:0] = {NbVid[7], NbVid[6:0]}. NbVid[7:0] specifies the Northbridge voltage. Writes outside the D18F5x17C[MaxVid, MinVid] range are ignored. Per <i>AMD Serial VID Interface 2.0 (SVI2) Specification</i> , effective VDDNB voltage should consider NbVid[7:0] and D18F5x188[NbOffsetTrim] .										
9:8	Reserved.										
7	NbDid: Northbridge divisor ID. Read-write. Specifies the Northbridge frequency divisor; see NbFid.										
6:1	NbFid: Northbridge frequency ID. Read-write. Specifies the Northbridge frequency multiplier. The NB COF is a function of NbFid and NbDid, and defined by NBCOF . NbFid and NbDid are not changed on a write if the value written results in a frequency greater than MSRC001_0071[MaxNb-Cof] . See 2.5.5 [P-state Bandwidth Requirements] .										
0	NbPstateEn: Northbridge P-state enable. Read-write. 1=The Northbridge P-state specified by this register is valid. 0=The Northbridge P-state specified by this register is not valid. This bit must be set to 1 in order for the Northbridge P-state specified by this register to be programmed in D18F5x170[NbPstateHi, NbPstateLo] .										

D18F5x170 Northbridge P-state Control

Updated-by-hardware. See also [2.5.4.1 \[NB P-states\]](#).

Bits	Description
31	MemPstateDis: memory P-state disable. IF (D18F3xE8[MemPstateCap]) THEN Read-write, updated-by-hardware. ELSE Read-only, updated-by-hardware. ENDIF. Reset: Value varies by product.. 1=Memory P-state transitions are disabled. The current P-state is not changed by programming this bit. The memory P-state will be forced to M0 on the next NB P-state transition. On processors where memory P-states are enabled, programming this bit may result in a violation of bandwidth requirements stated in 2.5.3.1.6 . Software must ensure that NB P-states which violate those requirements are forced disabled. 0=Memory P-state transitions are enabled if D18F2x90_dct[1:0][Dis-DllShutdownSR]=0 .
30	Reserved.

29:27	<p>NbPstateHiRes: NB P-state high residency timer. Read-write. Reset: 0. Specifies the minimum time the processor must spend in the high NB P-state before transitions to the low NB P-state are allowed. See 2.5.4.1 [NB P-states].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0us</td> <td>100b</td> <td>1ms</td> </tr> <tr> <td>001b</td> <td>10us</td> <td>101b</td> <td>5ms</td> </tr> <tr> <td>010b</td> <td>100us</td> <td>110b</td> <td>10ms</td> </tr> <tr> <td>011b</td> <td>500us</td> <td>111b</td> <td>50ms</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	0us	100b	1ms	001b	10us	101b	5ms	010b	100us	110b	10ms	011b	500us	111b	50ms
Bits	Description	Bits	Description																		
000b	0us	100b	1ms																		
001b	10us	101b	5ms																		
010b	100us	110b	10ms																		
011b	500us	111b	50ms																		
26:24	<p>NbPstateLoRes: NB P-state low residency timer. Read-write. Reset: 0. Specifies the minimum time the processor must spend in the low NB P-state before transitions to the high NB P-state are allowed. See: NbPstateHiRes.</p>																				
23	<p>NbPstateGnbSlowDis. Read-write. Reset: 0. Specifies whether NBP-state transitions take GPU activity into account. 0=Take GPU activity into account. 1=Ignore GPU activity .</p>																				
22:15	Reserved.																				
14	<p>SwNbPstateLoDis: software NB P-state low disable. IF (MSRC001_0071[NbPstateDis]) THEN Read-only. ELSE Read-write. ENDIF. BIOS: See 2.5.4.1.1.Reset: 0. 1=Transition to NbPstateHi and disable transitions to NbPstateLo.</p>																				
13	<p>NbPstateDisOnP0: NB P-state disable on P0. IF (MSRC001_0071[NbPstateDis]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0. 1=Transition to NbPstateHi and disable transitions to NbPstateLo if any compute unit is in P0 or a boosted P-state. This field uses software P-state numbering. See 2.5.3.1.2.1 [Software P-state Numbering].</p>																				
12	Reserved.																				
11:9	<p>NbPstateThreshold: NB P-state threshold. Read-write. Reset: Product-specific. BIOS: COUNT(D18F5x80[Enabled]). Specifies the minimum number of compute units that must be in a P-state with MSRC001_00[6B:64][NbPstate]=1 before transitions to lower performance NB P-states are allowed. See NbPstateLo and NbPstateHi.</p>																				
8	Reserved.																				
7:6	<p>NbPstateHi: NB P-state high. Read-write, updated-by-hardware. Cold reset: Product-specific. If NB P-states are enabled, this field specifies the NB P-state that is used when the number of compute units in a P-state with MSRC001_00[6B:64][NbPstate]=1 is less than NbPstateThreshold. This field must be programmed to the same value for all processors in the coherent fabric. This field is not changed on a write if the value written is greater than the NbPstateMaxVal value written or greater than the current NbPstateLo value. See also NbPstateDisOnP0, SwNbPstateLoDis, NbPstateLo, D18F5x174[NbPstateDis], and D18F5x1[6C:60][NbPstateEn].</p>																				
5	Reserved.																				
4:3	<p>NbPstateLo: NB P-state low. Read-write, updated-by-hardware. Cold reset: Product-specific. If NB P-states are enabled, this field specifies the NB P-state that is used when the number of compute units in a P-state with MSRC001_00[6B:64][NbPstate]=1 is greater than or equal to NbPstateThreshold. NbPstateLo must be greater than or equal to NbPstateHi. This field must be programmed to the same value for all processors in the coherent fabric. This field is not changed on a write if the value written is greater than the NbPstateMaxVal value written or less than the current NbPstateHi value. See also NbPstateDisOnP0, SwNbPstateLoDis, D18F5x174[NbPstateDis], and D18F5x1[6C:60][NbPstateEn].</p>																				
2	Reserved.																				

1:0	NbPstateMaxVal: NB P-state maximum value. Read-write. Cold reset: specified by the reset state of D18F5x1[6C:60][NbPstateEn] ; the cold reset value is the highest NB P-state number corresponding to the register in which NbPstateEn is set (e.g., if D18F5x160 and D18F5x164 have this bit set and the others do not, then NbPstateMaxVal=1; if NbPstateEn is only set in D18F5x160, then NbPstateMaxVal=0). This specifies the highest NB P-state value (lowest performance state) supported by the hardware.
-----	--

D18F5x174 Northbridge P-state Status

Bits	Description										
31:25	Reserved.										
24	CurMemPstate: current memory P-state. Read-only; updated-by-hardware. Reset: 0. Specifies the current memory P-state. 1=Memory P-state 1. 0=Memory P-state 0. See 2.5.7.1 [Memory P-states] .										
23	CurNbVid[7]: current northbridge voltage ID[7]. Read-only; updated-by-hardware. Reset: 0. MSRC001_0071[CurNbVid[7]] is an alias of D18F5x174[CurNbVid[7]] . .										
22:21	Reserved.										
20:19	CurNbPstate: current northbridge P-state. Read-only; updated-by-hardware. Cold reset: Product-specific. Provides the NB P-state that corresponds to the current frequency component of the NB. The value of this field is updated when the COF transitions to a new value associated with an NB P-state. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>NB P0</td> </tr> <tr> <td>01b</td> <td>NB P1</td> </tr> <tr> <td>10b</td> <td>NB P2</td> </tr> <tr> <td>11b</td> <td>NB P3</td> </tr> </tbody> </table>	Bits	Description	00b	NB P0	01b	NB P1	10b	NB P2	11b	NB P3
Bits	Description										
00b	NB P0										
01b	NB P1										
10b	NB P2										
11b	NB P3										
18:12	CurNbVid[6:0]: current northbridge voltage ID. MSRC001_0071[CurNbVid[6:0]] is an alias of D18F5x174[CurNbVid[6:0]] . Per <i>AMD Serial VID Interface 2.0 (SVI2) Specification</i> , effective VDDNB voltage should consider CurNbVid[7:0] and D18F5x188[NbOffsetTrim] .										
11:10	Reserved.										
9	CurNbDid: current northbridge divisor ID. Read-only, updated-by-hardware. Reset: 0.										
8:3	CurNbFid: current northbridge frequency ID. Read-only, updated-by-hardware. Reset: 0.										
2:1	StartupNbPstate: startup northbridge P-state number. Read-only. Cold reset: Product-specific. Specifies the cold reset VID, FID and DID for the Northbridge based on the NB P-state number selected.										
0	NbPstateDis: northbridge P-state disable. Read-only. Value: Product-specific. MSRC001_0071[NbPstateDis] is an alias of D18F5x174[NbPstateDis] .										

D18F5x178 Northbridge Fusion Configuration

Bits	Description
31:20	Reserved.
19	SwGfxDis. Read-write. Reset: 1. BIOS: IF (GpuEnabled) THEN 0. ELSE 1. ENDIF. This register is write once. 1=Hardware handshakes for NB P-state transitions and DRAM self-refresh entry are ignored. See 2.5.4.1.1 [NB P-state Transitions] and 2.5.7.2 [DRAM Self-Refresh] .

18	CstateFusionHsDis: C-state fusion handshake disable. Read-write. Reset: 0. BIOS: 1. 1=Ignore the FCH handshake response for PC6 transitions. 0=Use the FCH handshake response for PC6 entry. See 2.5.3.2.4.1 [FCH Messaging].
17	Dis2ndGnbAllowPsWait. Read-write. Reset: 0. BIOS: 1. 0=When doing an NB P-state transition wait for the display buffers to fill after quiescing cores. See 2.5.4.1.1 [NB P-state Transitions].
16:12	Reserved.
11	AllowSelfRefrS3Dis: allow self-refresh S3 disable. Read-write. Reset: 0. BIOS: 1. 1=The NB does not wait for handshake before placing DRAM into self-refresh (see 2.5.7.2 [DRAM Self-Refresh]) on S3 entry (see 2.5.8.1.1 [ACPI Suspend to RAM State (S3)]). 0=The NB waits for handshake before placing DRAM into self-refresh on S3 entry.
10	InbWakeS3Dis: InbWake S3 disable. Read-write. Reset: 0. BIOS: 1. 1= The NB does not wait for handshake before placing DRAM into self-refresh (see 2.5.7.2 [DRAM Self-Refresh]) on S3 entry (see 2.5.8.1.1 [ACPI Suspend to RAM State (S3)]). 0=The NB waits for handshake before placing DRAM into self-refresh on S3 entry.
9:4	Reserved.
3	CstateThreeWayHsEn: C-state three way handshake disable. Read-write. Reset: 0. 1=Enable the three way handshake with the FCH when entering a C-state. 0=Only a two way handshake with FCH is used. There is no message about the resulting package state sent to FCH. See 2.5.3.2.4.1 [FCH Messaging].
2	CstateFusionDis: C-state fusion disable. Read-write. Reset: 0. BIOS: 1. 1=All HALT or C-state requests are forwarded to the FCH. 0=HALT and C-state requests are forwarded to the FCH when each core has made a request. See 2.5.3.2.4.1 [FCH Messaging].
1:0	Reserved.

D18F5x17C Miscellaneous Voltages

Bits	Description
31	NbPsi0VidEn: Northbridge PSIO_L VID enable. Read-write. Reset: 0. BIOS: 2.5.1.3.1.1. This bit specifies how PSIO_L is controlled for VDDNB. See D18F3xA0[PsiVidEn] and 2.5.1.3.1 [PSIx_L Bit].
30:23	NbPsi0Vid[7:0]: Northbridge PSIO_L VID threshold. Read-write. Reset: 0. BIOS: 2.5.1.3.1.1. When enabled by NbPsi0VidEn, NbPsi0Vid specifies the threshold value of the VID code generated by the Northbridge, which in turn determines the state of PSIO_L. See D18F3xA0[PsiVid[6:0]] and 2.5.1.3.1 [PSIx_L Bit].
22:18	Reserved.
17:10	MinVid: minimum voltage. Read-only. Reset: Product-specific. Specifies the VID code corresponding to the minimum voltage (highest VID code) that the processor drives. 00h indicates that no minimum VID code is specified. See 2.5.1 [Processor Power Planes And Voltage Control].
9:8	Reserved.
7:0	MaxVid: maximum voltage. Read-only. Reset: Product-specific. Specifies the VID code corresponding to the maximum voltage (lowest VID code) that the processor drives. 00h indicates that no maximum VID code is specified. See 2.5.1 [Processor Power Planes And Voltage Control].

D18F5x188 Clock Power/Timing Control 5

Cold reset: 0000_000Eh. Updated-by-hardware. See the *AMD Serial VID Interface 2.0 (SVI2) Specification*.

Bits	Description
31:6	Reserved.
5	NbPsi1: Northbridge PS11_L. Read-write, updated-by-hardware. Specifies how PS11_L is controlled for VDDNB. 1=PS11_L is low. 0=PS11_L is high. See 2.5.1.3.1 [PSIx_L Bit] .
4:2	NbLoadLineTrim: Northbridge load line trim. Read-write. BIOS: D0F0xBC_xE0104184[SviLoadLineTrimVddNb] . See D18F5x12C[CoreLoadLineTrim] .
1:0	NbOffsetTrim: Northbridge offset trim. Read-write, updated-by-hardware. BIOS: D0F0xBC_xE0104184[SviLoadLineOffsetVddNb] . For encoding, see D18F5x12C[CoreOffsetTrim] .

D18F5x194 Name String Address Port

[D18F5x194](#) and [D18F5x198](#) provide BIOS with a read-only name string that may be copied to [MSRC001_00\[35:30\]](#) at warm reset. Each of [D18F5x198_x\[B:0\]](#) is read as follows:

1. Write [D18F5x194\[Index\]](#).
2. Read [D18F5x198](#).

Bits	Description						
31:4	Reserved.						
3:0	Index: name string register index. Read-write. Reset: 0. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Bh-0h</td> <td>Name String Registers. See D18F5x198_x[B:0].</td> </tr> <tr> <td>Fh-Ch</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	Bh-0h	Name String Registers. See D18F5x198_x[B:0] .	Fh-Ch	Reserved
Bits	Description						
Bh-0h	Name String Registers. See D18F5x198_x[B:0] .						
Fh-Ch	Reserved						

D18F5x198 Name String Data Port

See [D18F5x194](#) for register access information.

D18F5x198_x[B:0] Name String Data

Read-only. Value: Product-specific.

Bits	Description
31:24	NameStringByte3: name string ASCII character 3.
23:16	NameStringByte2: name string ASCII character 2.
15:8	NameStringByte1: name string ASCII character 1.
7:0	NameStringByte0: name string ASCII character 0.

3.14 GPU Memory Mapped Registers

GMMx63C CFG_VOLTAGE_CONTROL

See 2.5.1.4.2.1 [Software-Initiated NB Voltage Transitions].

Bits	Description
31:16	Reserved.
15:8	VoltageLevel. Read-write. Reset: 0. Specifies the voltage requested when software toggles GMMx63C[VoltageChangeReq] . Values outside the D18F5x17C[MaxVid, MinVid] range are invalid and result in undefined behavior.
7:3	Reserved.
2	VoltageChangeReq. Read-write. Reset: 0. Software toggles this field to make VDDNB voltage requests.
1	VoltageChangeEn. Read-write. Reset: 0. Specifies whether changes to GMMx63C[VoltageChangeReq] causes voltage change requests. 1=Requests occur. 0=Requests do not occur.
0	VoltageForceEn. Read-write. Reset: 0. If GMMx63C[VoltageChangeEn] ==1, this field specifies whether changes to GMMx63C[VoltageChangeReq] caused forced voltage changes. 1=Voltage changes are forced. 0=Voltage changes are not forced.

GMMx640 CFG_VOLTAGE_STATUS

See 2.5.1.4.2.1 [Software-Initiated NB Voltage Transitions].

Bits	Description
31:9	Reserved.
8:1	CurrentVoltageLevel. Read-only. Reset: 0. Specifies the current voltage level requested by GMMx63C . See GMMx63C[VoltageLevel] . To determine the current voltage level, software must poll on this field until two consecutive reads return the same value.
0	VoltageChangeAck. Read-only. Reset: 0. Specifies whether the voltage change requested by GMMx63C[VoltageChangeReq] is complete.

GMMx770 CG Voltage Control

Reset: 0000_0006h. See 2.5.1.4.2 [Software-Initiated Voltage Transitions].

Bits	Description
31:11	Reserved.
10	VoltageForceEn. Read-write. See GMMx63C[VoltageForceEn] .
9	VoltageChangeEn. Read-write. See GMMx63C[VoltageChangeEn] .
8:1	VoltageLevel. Read-write. See GMMx63C[VoltageLevel] .
0	VoltageChangeReq. Read-write. See GMMx63C[VoltageChangeReq] .

GMMx774 CG Voltage Status

Reset: 0000_0000h. See 2.5.1.4.2 [Software-Initiated Voltage Transitions].

Bits	Description
31:9	Reserved.
8:1	CurrentVoltageLevel . Read-only. See GMMx640[CurrentVoltageLevel] .
0	VoltageChangeAck . Read-only. See GMMx640[VoltageChangeAck] .

GMMx7A0 LCLK_DEEP_SLEEP_CNTL

Bits	Description						
31	EnableDs . Read-write. Reset: 0. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Deep Sleep Disabled.</td> </tr> <tr> <td>1</td> <td>Deep Sleep Enabled.</td> </tr> </tbody> </table>	Bits	Definition	0	Deep Sleep Disabled.	1	Deep Sleep Enabled.
Bits	Definition						
0	Deep Sleep Disabled.						
1	Deep Sleep Enabled.						
30	DmaactiveMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Allow deep sleep on DMAACTIVE_L</td> </tr> <tr> <td>1</td> <td>DMAACTIVE_L has to be deasserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Allow deep sleep on DMAACTIVE_L	1	DMAACTIVE_L has to be deasserted to enter deep sleep
Bits	Definition						
0	Allow deep sleep on DMAACTIVE_L						
1	DMAACTIVE_L has to be deasserted to enter deep sleep						
29	OnOutbWakeAckMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Allow deep sleep on OnOutbWakeAck</td> </tr> <tr> <td>1</td> <td>OnOutbWakeAck has to be deasserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Allow deep sleep on OnOutbWakeAck	1	OnOutbWakeAck has to be deasserted to enter deep sleep
Bits	Definition						
0	Allow deep sleep on OnOutbWakeAck						
1	OnOutbWakeAck has to be deasserted to enter deep sleep						
28	OnOutbWakeMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Allow deep sleep on OnOutbWake</td> </tr> <tr> <td>1</td> <td>OnOutbWake has to be deasserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Allow deep sleep on OnOutbWake	1	OnOutbWake has to be deasserted to enter deep sleep
Bits	Definition						
0	Allow deep sleep on OnOutbWake						
1	OnOutbWake has to be deasserted to enter deep sleep						
27	OnInbWakeAckMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Allow deep sleep on OnInbWakeAck</td> </tr> <tr> <td>1</td> <td>OnInbWakeAck has to be deasserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Allow deep sleep on OnInbWakeAck	1	OnInbWakeAck has to be deasserted to enter deep sleep
Bits	Definition						
0	Allow deep sleep on OnInbWakeAck						
1	OnInbWakeAck has to be deasserted to enter deep sleep						
26	OnInbWakeMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Allow deep sleep on OnInbWake</td> </tr> <tr> <td>1</td> <td>OnInbWake has to be deasserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Allow deep sleep on OnInbWake	1	OnInbWake has to be deasserted to enter deep sleep
Bits	Definition						
0	Allow deep sleep on OnInbWake						
1	OnInbWake has to be deasserted to enter deep sleep						
25	OrbIdleMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Do not wait for ORB_SMU_idle to enter deep sleep</td> </tr> <tr> <td>1</td> <td>ORB_SMU_idle has to be asserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Do not wait for ORB_SMU_idle to enter deep sleep	1	ORB_SMU_idle has to be asserted to enter deep sleep
Bits	Definition						
0	Do not wait for ORB_SMU_idle to enter deep sleep						
1	ORB_SMU_idle has to be asserted to enter deep sleep						
24	L2imuIdleMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Do not wait for L2IMU_SMU_idle to enter deep sleep</td> </tr> <tr> <td>1</td> <td>L2IMU_SMU_idle has to be asserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Do not wait for L2IMU_SMU_idle to enter deep sleep	1	L2IMU_SMU_idle has to be asserted to enter deep sleep
Bits	Definition						
0	Do not wait for L2IMU_SMU_idle to enter deep sleep						
1	L2IMU_SMU_idle has to be asserted to enter deep sleep						
23	L1imuintgenIdleMask . Read-write. Reset: 1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Do not wait for L1IMUINTGEN_SMU_idle to enter deep sleep</td> </tr> <tr> <td>1</td> <td>L1IMUINTGEN_SMU_idle has to be asserted to enter deep sleep</td> </tr> </tbody> </table>	Bits	Definition	0	Do not wait for L1IMUINTGEN_SMU_idle to enter deep sleep	1	L1IMUINTGEN_SMU_idle has to be asserted to enter deep sleep
Bits	Definition						
0	Do not wait for L1IMUINTGEN_SMU_idle to enter deep sleep						
1	L1IMUINTGEN_SMU_idle has to be asserted to enter deep sleep						

22	L1imubifIdleMask. Read-write. Reset: 1.
	<u>Bits</u> <u>Definition</u>
	0 Do not wait for L1IMUBIF_SMU_idle to enter deep sleep 1 L1IMUBIF_SMU_idle has to be asserted to enter deep sleep
21	L1imugppsbIdleMask. Read-write. Reset: 1.
	<u>Bits</u> <u>Definition</u>
	0 Do not wait for L1IMUGPPSB_SMU_idle to enter deep sleep 1 L1IMUGPPSB_SMU_idle has to be asserted to enter deep sleep
20	L1imugfxIdleMask. Read-write. Reset: 1.
	<u>Bits</u> <u>Definition</u>
	0 Do not wait for L1IMUGFX_SMU_idle to enter deep sleep 1 L1IMUGFX_SMU_idle has to be asserted to enter deep sleep
19	PcieLclkIdle2Mask. Read-write. Reset: 1.
	<u>Bits</u> <u>Definition</u>
	0 Do not wait for PCIE_SMU_lclk_idle2 to enter deep sleep 1 PCIE_SMU_lclk_idle2 has to be asserted to enter deep sleep
18	PcieLclkIdle1Mask. Read-write. Reset: 1.
	<u>Bits</u> <u>Definition</u>
	0 Do not wait for PCIE_SMU_lclk_idle1 to enter deep sleep 1 PCIE_SMU_lclk_idle1 has to be asserted to enter deep sleep
17	SmuBusyMask. Read-write. Reset: 1.
	<u>Bits</u> <u>Definition</u>
	0 Do not wait for smu stop to enter Deep Sleep 1 SMU stop has to be asserted to enter Deep Sleep.
16	SclkRunningMask. Read-write. Reset: 1.
	<u>Bits</u> <u>Definition</u>
	0 Do not wait for GFX Idle to enter Deep Sleep 1 GFX should be idle to enter Deep sleep.
15:4	Hysteresis. Read-write. Reset: 0.
3	RampDis. Read-write. Reset: 0.
	<u>Bits</u> <u>Definition</u>
	0 Clock will be ramped up and down. 1 Clock ramping disabled.
2:0	DivId. Read-write. Reset: 5h.
	<u>Bits</u> <u>Definition</u>
	000b Clock OFF
	001b Divide by 2.
	010b Divide by 4.
	011b Divide by 8.
	100b Divide by 16.
	101b Divide by 32.
110b Reserved.	
111b Reserved.	

3.15 IOMMU Memory Mapped Registers

See 3.1 [Register Descriptions and Mnemonics] for a description of the register naming convention. See 2.12.1 [IOMMU Configuration Space].

IOMMUx00 Device Table Base Address Low

Bits	Description
31:12	DevTblBase[31:12]: device table base address bits[31:12]. Read-write. Reset: 0. DevTblBase[51:12] = {IOMMUx04[DevTblBase[51:32], DevTblBase[31:12]}. DevTblBase[51:12] specifies the 4Kbyte-aligned base address of the first level device table.
11:9	Reserved1. Reserved for future use. This register controls no hardware
8:0	DevTblSize: device table size. Read-write. Reset: 0. This field contains 1 less than the length of the device table, in multiples of 4K bytes. A minimum size of 0 corresponds to a 4K byte device table and a maximum size of 1FFh corresponds to a 2M byte device table.

IOMMUx04 Device Table Base Address High

Bits	Description
31:20	Reserved.
19:0	DevTblBase[51:32]: device table base address bits[51:32]. See: IOMMUx00[DevTblBase[31:12]].

IOMMUx08 Command Buffer Base Address Low

Bits	Description
31:12	ComBase[31:12]: command buffer base address bits[31:12]. Read-write. Reset: 0. ComBase[51:12] = {IOMMUx0C[ComBase[51:32], ComBase[31:12]}. ComBase[51:12] specifies the 4Kbyte-aligned base address of the command buffer.
11:0	Reserved.

IOMMUx0C Command Buffer Base Address High

Bits	Description						
31:28	Reserved.						
27:24	ComLen: command buffer length. Read-write. Reset: 8h. Specifies the length of the command buffer in power of 2 increments. The minimum size is 256 entries (4K bytes); values less than 8h are reserved.						
	<table border="0"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>7h-0h</td> <td>Reserved.</td> </tr> <tr> <td>Fh-8h</td> <td>2^{ComLen} entries (2^{ComLen}*16 bytes).</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	7h-0h	Reserved.	Fh-8h	2 ^{ComLen} entries (2 ^{ComLen} *16 bytes).
<u>Bits</u>	<u>Description</u>						
7h-0h	Reserved.						
Fh-8h	2 ^{ComLen} entries (2 ^{ComLen} *16 bytes).						

23:20	Reserved.
19:0	ComBase[51:32]: command buffer base address bits[51:32]. See: IOMMUx08[Com-Base[31:12]] .

IOMMUx10 Event Log Base Address Low

Bits	Description
31:12	EventBase[31:12]: event log base address bits[31:12]. Read-write. Reset: 0. EventBase[51:12] = { IOMMUx14[EventBase[51:32] , EventBase[31:12]}. EventBase[51:12] specifies the 4K-byte aligned base address of the event log.
11:0	Reserved.

IOMMUx14 Event Log Base Address High

Bits	Description						
31:28	Reserved.						
27:24	EventLen: event log length. Read-write. Reset: 8h. Specifies the length of the event log in power of 2 increments. The minimum size is 256 entries (4K bytes); values less than 8h are reserved. <table border="1" data-bbox="272 934 1055 1039"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7h-0h</td> <td>Reserved.</td> </tr> <tr> <td>Fh-8h</td> <td>2^{EventLen} entries (2^{EventLen}*16 bytes).</td> </tr> </tbody> </table>	Bits	Description	7h-0h	Reserved.	Fh-8h	2 ^{EventLen} entries (2 ^{EventLen} *16 bytes).
Bits	Description						
7h-0h	Reserved.						
Fh-8h	2 ^{EventLen} entries (2 ^{EventLen} *16 bytes).						
23:20	Reserved.						
19:0	EventBase[51:32]: event log base address bits [51:32]. See: IOMMUx10[EventBase[31:12]] .						

IOMMUx18 Control Low

Bits	Description
31:22	Reserved.
21:18	Tlpt. Read-write. Reset: 0. Tlpt contains the 4-bit value matched to the PCIe TLP Type field when the PCIe TLP Fmt value indicates the field carries a prefix.
17	GaEn. Read-write. Reset: 0. Guest APIC enable. 1 = loose. 0 = prohibited
16	GtEn. Read-write. Reset: 0. 1=Guest translation may be enabled for a peripheral by programming DTE[GV]. This bit must be programmed to zero when IOMMUx30[GtSup]=0 .
15	PprEn. Read-write. Reset: 0. 1=Peripheral page service requests are processed. 0=Peripheral page service requests are treated as invalid device requests. This bit must be programmed to zero when IOMMUx30[PprSup]=0 .
14	PprIntEn. Read-write. Reset: 0. 1=An interrupt is generated when IOMMUx2020[PprInt]=1 or IOMMUx2020[PprOverflow]=1 . The interrupt vector used is indicated in D0F2x50[IommuM-siNumPpr] . This bit must be programmed to zero when IOMMUx30[PprSup]=0 .

13	<p>PprLogEn. Read-write. Reset: 0. 1=Peripheral page service request events are written to the peripheral page service request log when IommuEn=1. 0=Peripheral page service request logging is not enabled. Peripheral page service requests are discarded when PprLogEn=0 or IOMMUx30[PprSup]=0. When IommuEn=1 and software sets PprLogEn, the IOMMU clears IOMMUx2020[PprOverflow] and sets IOMMUx2020[PprRun]. The IOMMU can then write new entries to the event log if there are usable entries available. Software can read IOMMUx2020[PprRun] to determine the status of the peripheral page service request log. Note the peripheral page service request and event logs are independent. IOMMUx38, IOMMUx2030, and IOMMUx2038 must be programmed prior to setting PprLogEn.</p>																
12	<p>CmdBufEn. Read-write. Reset: 0. 1=Start or restart command buffer processing. When CmdBufEn=1 and IommuEn=1, the IOMMU starts fetching commands and sets IOMMUx2020[CmdBufRun]. 0=Halt command buffer processing. Writing a 0 to this bit causes the IOMMU to cease fetching new commands although commands previously fetched are completed. The IOMMU stops fetching commands upon reset and after errors. See IOMMUx2020[CmdBufRun]. Writing of event log entries is independently controlled by EventLogEn. IOMMUx08, IOMMUx0C, IOMMUx2000, and IOMMUx2008 must be programmed prior to setting CmdBufEn.</p>																
11	<p>Isoc. Read-write. Reset: 0. This bit controls the state of the isochronous bit in the HyperTransport read request packet when the IOMMU issues I/O page table reads and device table reads on the HyperTransport link. 1=Request packet to use isochronous channel. 0=Request packet to use standard channel. This bit must be set if IommuEn==1.</p>																
10	<p>Coherent. Read-write. Reset: 1. This bit controls the state of the coherent bit in the HyperTransport read request packet when the IOMMU issues device table reads on the HyperTransport link. 1=Device table requests are snooped by the processor. 0=Device table requests are not snooped by the processor.</p>																
9	<p>ResPassPw. Read-write. Reset: 0. This bit controls the state of the ResPassPW bit in the HyperTransport read request packet when the IOMMU issues I/O page table reads and device table reads on the HyperTransport link. 1=Response may pass posted requests. 0=Response may not pass posted requests.</p>																
8	<p>PassPw. Read-write. Reset: 0. This bit controls the state of the PassPW bit in the HyperTransport read request packet when the IOMMU issues I/O page table reads and device table reads on the HyperTransport link.. 1=Request packet may pass posted requests. 0=Request packet may not pass posted requests.</p>																
7:5	<p>InvTimeout. Read-write. Reset: 0. This field specifies the invalidation timeout for IOTLB invalidation requests.</p> <table border="1" data-bbox="280 1430 670 1703"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>No timeout.</td> </tr> <tr> <td>001b</td> <td>1 ms.</td> </tr> <tr> <td>010b</td> <td>10 ms.</td> </tr> <tr> <td>011b</td> <td>100 ms.</td> </tr> <tr> <td>100b</td> <td>1 sec.</td> </tr> <tr> <td>101b</td> <td>10 sec.</td> </tr> <tr> <td>111b-110b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	000b	No timeout.	001b	1 ms.	010b	10 ms.	011b	100 ms.	100b	1 sec.	101b	10 sec.	111b-110b	Reserved
<u>Bits</u>	<u>Description</u>																
000b	No timeout.																
001b	1 ms.																
010b	10 ms.																
011b	100 ms.																
100b	1 sec.																
101b	10 sec.																
111b-110b	Reserved																
4	<p>ComWaitInten. Read-write. Reset: 0. 1=An interrupt is generated when IOMMUx2020[ComWaitInt]=1.</p>																
3	<p>EventIntEn. Read-write. Reset: 0. 1=An interrupt is generated when IOMMUx2020[EventLogInt]=1 or IOMMUx2020[EventOverflow]=1.</p>																

2	EventLogEn. Read-write. Reset: 0. 1=All events detected are written to the event log when IommuEn=1. 0=Event logging is not enabled. Events are discarded when the event log is not enabled. When IommuEn=1 and software sets EventLogEn, the IOMMU clears IOMMUx20[EventOverflow] and sets IOMMUx20[EventLogRun] . IOMMUx10 , IOMMUx14 , IOMMUx2010 , and IOMMUx2018 must be programmed prior to setting EventLogEn.
1	HtTunEn. Read-write. Reset: 0. 1=Upstream traffic received by the HyperTransport tunnel is translated by the IOMMU. 0=Upstream traffic received by the HyperTransport tunnel is not translated by the IOMMU. The IOMMU ignores the state of this bit while IommuEn=0. See DOF2x40[IommuHtTunnelSup] .
0	IommuEn. Read-write. Reset: 0. 1=IOMMU enabled. All upstream transactions are translated by the IOMMU. IOMMUx00 [Device Table Base Address Low] and IOMMUx04 [Device Table Base Address High] must be configured by software before setting this bit. 0=IOMMU is disabled and no upstream transactions are translated or remapped by the IOMMU. When disabled, the IOMMU does not read any commands or create any event log entries.

IOMMUx20 Exclusion Range Base Low

Bits	Description
31:12	ExclBase[31:12]: exclusion range base address bits[31:12]. Read-write. Reset: 0. ExclBase[51:12] = { IOMMUx20[ExclBase[51:32]] , ExclBase[31:12]}. Specifies the 4Kbyte-aligned base address of the exclusion range.
11:2	Reserved.
1	ExAllow: exclusion allow. Read-write. Reset: 0. 1=All accesses to the exclusion range are forwarded untranslated. 0=The EX bit in the device table entry specifies if accesses to the exclusion range are translated.
0	ExEn: exclusion enable. Read-write. Reset: 0. 1=The exclusion range is enabled.

IOMMUx24 Exclusion Range Base High

Bits	Description
31:20	Reserved.
19:0	ExclBase[51:32]: exclusion range base address bits[51:32]. See: IOMMUx20[ExclBase[31:12]] .

IOMMUx28 Exclusion Range Limit Low

Bits	Description
31:12	ExclLimit[31:12]: exclusion range limit address bits[31:12]. Read-write. Reset: 0. ExclLimit[51:12] = { IOMMUx2C[ExclLimit[51:32]] , ExclLimit[31:12]}. ExclLimit[51:12] specifies the 4Kbyte-aligned limit address of the exclusion range.
11:0	Reserved.

IOMMUx2C Exclusion Range Limit High

Bits	Description
31:20	Reserved.
19:0	ExclLimitHi . See: IOMMUx28 [ExclLimit[31:12]].

IOMMUx30 Extended Feature Low

Bits	Description										
31:15	Reserved.										
14:13	GlXSup . Read-only. Reset: 00b. 01b=Two-level GCR3 base address table is supported in hardware. 0=single-level GCR3 base table address translation is supported. 00b=GLX in the DTE is ignored and the IOMMU performs only single-level guest CR3 lookups. This value is not meaningful when GtSup=0.										
12	Reserved.										
11:10	HATS: host address translation size . Read-only. Reset: 10b. The maximum number of host address translation levels supported. This value is not meaningful when GtSup=0. IOMMU behaviour is undefined if Next Level in a page directory entry exceeds the limit set by HATS. <table border="1" data-bbox="284 940 673 1108"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>4 levels.</td> </tr> <tr> <td>01b</td> <td>5 levels.</td> </tr> <tr> <td>10b</td> <td>6 levels.</td> </tr> <tr> <td>11b</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	00b	4 levels.	01b	5 levels.	10b	6 levels.	11b	Reserved.
Bits	Description										
00b	4 levels.										
01b	5 levels.										
10b	6 levels.										
11b	Reserved.										
9	PcSup: performance counters supported . Read-only. Reset: 1. 1=performance counters are supported.										
8	HeSup: hardware error registers supported . Read-only. Reset: 0. 0=Hardware error registers do not report error information.. 1=Error information is reported in hardware error registers.										
7	GaSup: guest APIC supported . Read-only. Reset: 0. 1=Guest APIC supported.										
6	IaSup: INVALIDATE_IOMMU_ALL supported . Read-only. Reset: 1. 1=The INVALIDATE_IOMMU_ALL command is supported.										
5	Reserved.										
4	GtSup: guest translation supported . Read-only. Reset: 1. 1=Guest address translation is supported. 0=Only nested address translation is supported. When GtSup=0, the following values in the DTE must be zero: GV, GLX and GCR3 Table Root Pointer. See IOMMUx18 [GtEn].										
3	NxSup: no execute supported . Read-only. Reset: 0. 1=No-execute protection is supported.. 0=No-execute protection is not supported.										
2	XtSup: x2 apic supported . Read-only. Reset: 0. 1=The interrupt remapping table is expanded to support x2APIC interrupt information.. 0=x2APIC support is disabled.										
1	PprSup: peripheral page service request supported . Read-only. Reset: 1. 1=Indicates that IOMMU handles page service request events from peripherals, the IOMMU supports the page service request queue, and that the second IOMMU interrupt can be used to signal peripheral page service request events.										
0	PrefSup: prefetch support . Read-only. Reset: 1. 1=Indicates that IOMMU will accept PREFETCH_IOMMU_PAGES commands.										

IOMMUx34 Extended Feature High

Bits	Description						
31:4	Reserved.						
3:0	PasMax: PASID maximum. Read-only. Reset: 8h. This specifies the maximum PASID value supported. This field is not meaningful when IOMMUx34[GsSup] =0. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>2h-0h</td> <td>Reserved.</td> </tr> <tr> <td>Fh-3h</td> <td>$2^{(PasMax+1)}-1$.</td> </tr> </tbody> </table>	Bits	Description	2h-0h	Reserved.	Fh-3h	$2^{(PasMax+1)}-1$.
Bits	Description						
2h-0h	Reserved.						
Fh-3h	$2^{(PasMax+1)}-1$.						

IOMMUx38 PPR Log Base Address Low

Bits	Description
31:12	PprBase[31:12]: ppr base address bits[31:12]. Read-write. Reset: 0. $PprBase[51:12] = \{IOMMUx3C[PprBase[51:32]], PprBase[31:12]\}$. $PprBase[51:12]$ specifies the 4Kbyte-aligned base address of the PPR log.
11:0	Reserved.

IOMMUx3C PPR Log Base Address High

Bits	Description						
31:28	Reserved.						
27:24	PprLen: ppr length. Read-write. Reset: 8h. Specifies the length of the PPR log in power of two increments. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7h-0h</td> <td>Reserved.</td> </tr> <tr> <td>Fh-8h</td> <td>2^{PprLen} entries ($2^{PprLen} * 16$ bytes).</td> </tr> </tbody> </table>	Bits	Description	7h-0h	Reserved.	Fh-8h	2^{PprLen} entries ($2^{PprLen} * 16$ bytes).
Bits	Description						
7h-0h	Reserved.						
Fh-8h	2^{PprLen} entries ($2^{PprLen} * 16$ bytes).						
23:20	Reserved.						
19:0	PprBase[51:32]: ppr base address bits[51:32]. See: IOMMUx38[31:12] .						

IOMMUx40 Hardware Error Upper Low

Bits	Description
31:0	FirstEvCode[31:0]: first event code bits[31:0]. Read-write. Reset: 0. $FirstEvCode[59:0] = \{IOMMUx44[FirstEvCode[59:32]], FirstEvCode[31:0]\}$. IOMMUx44[EvCode] and $FirstEvCode[59:0]$ specify the upper 64 bits of the most recent hardware error detected by the IOMMU.

IOMMUx44 Hardware Error Upper High

Bits	Description
31:28	EvCode: event code. Read-write. Reset: 0. Event code for the type of error logged.
27:0	FirstEvCode[59:32]: first event code bits[59:32]. See: IOMMUx40[FirstEvCode[31:0]] .

IOMMUx48 Hardware Error Lower Low

Bits	Description
31:0	SecondEvCode[31:0]: second event code bits[31:0] . Read-write. Reset: 0. SecondEvCode[63:0] = {IOMMUx4C[SecondEvCode[63:32]], SecondEvCode[31:0]}. SecondEvCode[63:0] specifies the lower 64 bits of the most recent hardware error detected by IOMMU.

IOMMUx4C Hardware Error Lower High

Bits	Description
31:0	SecondEvCode[63:32]: second event code bits[63:32] . See: IOMMUx48[SecondEvCode[31:0]].

IOMMUx50 Hardware Error Status

Bits	Description
31:2	Reserved.
1	HEO: hardware error overflow . Read-write. Reset: 0. Defines the contents of the IOMMU hardware error registers as having been overwritten.. 0=not overwritten.. 1=contents overwritten by new information.. HEO is not meaningful when HEV=0.
0	HEV: hardware error valid . Read-write. Reset: 0. 1=Contents of the IOMMU hardware error registers are valid.

IOMMUx2000 Command Buffer Head Pointer

Bits	Description
31:19	Reserved.
18:4	CmdHdptr: command buffer head pointer . Read-write. Reset: 0. Specifies the 128-bit aligned offset from the command buffer base address register of the next command to be fetched by the IOMMU. The IOMMU increments this register, rolling over to zero at the end of the buffer, after fetching and validating the command in the command buffer. After incrementing this register, the IOMMU cannot re-fetch the command from the buffer. If this register is written by software while IOMMUx2020[CmdBufRun]=1, the IOMMU behavior is undefined. If this register is set by software to a value outside the length specified by IOMMUx08[ComLen], the IOMMU behavior is undefined.
3:0	Reserved.

IOMMUx2008 Command Buffer Tail Pointer

Bits	Description
31:19	Reserved.

18:4	CmdTailptr: command buffer tail pointer. Read-write. Reset: 0. Specifies the 128-bit aligned offset from the command buffer base address register of the next command to be written by the software. Software must increment this field, rolling over to zero at the end of the buffer, after writing a command to the command buffer. If software advances the tail pointer equal to or beyond the head pointer after adding one or more commands to the buffer, the IOMMU behavior is undefined. If software sets the command buffer tail pointer to an offset beyond the length of the command buffer, the IOMMU behavior is undefined.
3:0	Reserved.

IOMMUx2010 Event Log Head Pointer

Bits	Description
31:19	Reserved.
18:4	EventHdptr: event log head pointer. Read-write. Reset: 0. Specifies the 128 bit aligned offset from the event log base address register that will be read next by software. Software must increment this field, rolling over at the end of the buffer, after reading an event from the event log. If software advances the head pointer beyond the tail pointer, the IOMMU behavior is undefined. If software sets the event log head pointer to an offset beyond the length of the event log, the IOMMU behavior is undefined.
3:0	Reserved.

IOMMUx2018 Event Log Tail Pointer

Bits	Description
31:19	Reserved.
18:4	EventTailptr: event log tail pointer. Read-write. Reset: 0. Specifies the 128-bit aligned offset from the event log base address register that will be written next by the IOMMU when an event is detected. The IOMMU increments this register, rolling over at the end of the buffer, after writing an event to the event log. If this register is written while IOMMUx2020[EventLogRun]=1 , the IOMMU behavior is undefined. If this register is set by software to a value outside the length specified by IOMMUx10[EventLen] , the IOMMU behavior is undefined.
3:0	Reserved.

IOMMUx2020 Status

Bits	Description
31:8	Reserved1. Reserved for future use. This register controls no hardware
7	PprRun: peripheral page service request running. Read-only. Reset: 0. 1=PPR requests are logged as they occur. 0=PPR requests are discarded without logging.. When PprOverflow=1, the IOMMU does not write new PPR log entries even when PprRun=1. When halted, PPR request logging is restarted by using IOMMUx18[PprLogEn] .
6	PprInt: peripheral page service request interrupt. Read-write; Write-1-to-clear. Reset: 0. 1=PPR request entry written to the PPR log by the IOMMU. 0=No PPR entry written to the PPR log by the IOMMU. See IOMMUx18[PprIntEn] .

5	PprOverflow: peripheral page service request overflow. Read-write; Write-1-to-clear. Reset: 0. 1=IOMMU PPR log overflow has occurred. This bit is set when a new peripheral page service request is to be written to the PPR log and there is no usable entry in the PPR log, causing the new information to be discarded. No new PPR log entries are written while this bit is set. See IOMMUx18[PprIntEn] .
4	CmdBufRun: command buffer running. Read-only. Reset: 0. 1=Commands may be fetched from the command buffer. 0=IOMMU has stopped fetching new commands. The IOMMU freezes command processing after <code>COMMAND_HARDWARE_ERROR</code> or <code>ILLEGAL_COMMAND_ERROR</code> errors. When frozen, command fetching is restarted by using IOMMUx18[CmdBufEn] .
3	EventLogRun: event log running. Read-only. Reset: 0. 1=Events are logged as they occur.. 0=Event reports are discarded without logging. When <code>EventOverflow=1</code> , the IOMMU does not write new event log entries even when <code>EventLogRun=1</code> . When halted, event logging is restarted by using IOMMUx18[EventLogEn] .
2	ComWaitInt: completion wait interrupt. Read-write; Write-1-to-clear. Reset: 0. 1= <code>COMPLETION_WAIT</code> command completed. This bit is only set if the <code>i</code> bit is set in the <code>COMPLETION_WAIT</code> command. See IOMMUx18[ComWaitIntEn] .
1	EventLogInt: event log interrupt. Read-write; Write-1-to-clear. Reset: 0. 1=Event entry written to the event log by the IOMMU. See IOMMUx18[EventIntEn] .
0	EventOverflow. Read-write; Write-1-to-clear. Reset: 0. 1=IOMMU event log overflow has occurred. This bit is set when a new event is to be written to the event log and there is no usable entry in the event log, causing the new event information to be discarded. No new event log entries are written while this bit is set. See IOMMUx18[EventIntEn] .

IOMMUx2030 PPR Log Head Pointer

Bits	Description
31:19	Reserved.
18:4	PprHdptr: ppr head pointer. Read-write. Reset: 0. Specifies the 128-bit aligned offset from the PPR log base address register that will be read next by software. Software must increment this field, rolling over at the end of the buffer, after reading a PPR request entry from the PPR event log. If software advances the head pointer beyond the tail pointer, the IOMMU behavior is undefined. If software sets the PPR log head pointer to an offset beyond the length of the PPR log, the IOMMU behavior is undefined.
3:0	Reserved.

IOMMUx2038 PPR Log Tail Pointer

Bits	Description
31:19	Reserved.

18:4	PprTailptr. Read-write. Reset: 0. Specifies the 128-bit aligned offset from the PPR log base address register that will be written next by the IOMMU when a PPR request is detected. The IOMMU increments this register, rolling over at the end of the buffer, after writing a PPR request to the PPR log. If this register is written while <code>IOMMUx2020[PprRun]=1</code> , the IOMMU behavior is undefined. If software sets the PPR log tail pointer to an offset beyond the length of the PPR log, defined by <code>IOMMUx3C[PprLen]</code> , the IOMMU behavior is undefined.
3:0	Reserved.

IOMMUx4000 Counter Configuration

Bits	Description								
31:18	Reserved.								
17:12	NCounterBanks: number of counter banks. Read-only. Reset: 2h. The number of counter banks supported by the IOMMU. Each bank contains two or more counter and control registers as specified by NCounter. For each counter bank, a corresponding control bit is in <code>IOMMUx4008</code> , <code>IOMMUx4010</code> , and <code>IOMMUx4018</code> . Each supported event counter bank is in a distinct, consecutive 4K byte page. <table border="0"> <tr> <td><u>Bits</u></td> <td><u>Description</u></td> </tr> <tr> <td>00h</td> <td>No counter banks supported.</td> </tr> <tr> <td>3Fh-01h</td> <td>NCounterBanks event counter banks are supported.</td> </tr> </table> Note: IOMMU event counter banks are numbered starting with 0.	<u>Bits</u>	<u>Description</u>	00h	No counter banks supported.	3Fh-01h	NCounterBanks event counter banks are supported.		
<u>Bits</u>	<u>Description</u>								
00h	No counter banks supported.								
3Fh-01h	NCounterBanks event counter banks are supported.								
11	Reserved.								
10:7	NCounter: number of counters per bank. Read-only. Reset: 4h. Reports the number of individual counters in each IOMMU counter bank. Each counter bank contains the same number of counters. <table border="0"> <tr> <td><u>Bits</u></td> <td><u>Description</u></td> </tr> <tr> <td>0h</td> <td>No counters supported.</td> </tr> <tr> <td>1h</td> <td>Reserved.</td> </tr> <tr> <td>Fh-2h</td> <td>NCounter counters in each bank.</td> </tr> </table>	<u>Bits</u>	<u>Description</u>	0h	No counters supported.	1h	Reserved.	Fh-2h	NCounter counters in each bank.
<u>Bits</u>	<u>Description</u>								
0h	No counters supported.								
1h	Reserved.								
Fh-2h	NCounter counters in each bank.								
6:0	Reserved.								

IOMMUx4008 Counter PASID Bank Lock Low

Bits	Description
31:0	PasidLock[31:0]: pasid lock enable bits[31:0]. Read-write. Reset: 0. <code>PasidLock[63:0] = {IOMMUx400C[PasidLock[63:32]], PasidLock[31:0]}</code> . For each bit in <code>PasidLock[63:0]</code> , if the bit is set then writes to the corresponding bank in <code>IOMMUx4[1,0][3:0]10</code> and <code>IOMMUx4[1,0][3:0]14</code> are ignored. Bit positions above the value reported in <code>IOMMUx4000[NCounterBanks]</code> are ignored when written and return zero when read.

IOMMUx400C Counter PASID Bank Lock High

Bits	Description
31:0	PasidLock[63:32]. See: <code>IOMMUx4008[PasidLock[31:0]]</code> .

IOMMUx4010 Domain Bank Lock Low

Bits	Description
31:0	DomainLock[31:0]: domain lock enable bits[31:0] . Read-write. Reset: 0. DomainLock[63:0] = {IOMMUx4014[DomainLock[63:32]], DomainLock[31:0]}. For each bit in DomainLock[63:0], if the bit is set then writes to the corresponding bank in IOMMUx4[1,0][3:0]18 and IOMMUx4[1,0][3:0]1C are ignored. Bit positions above the value reported in IOMMUx4000[NCounterBanks] are ignored when written and return zero when read.

IOMMUx4014 Domain Bank Lock High

Bits	Description
31:0	DomainLock[63:32] . See: IOMMUx4010[DomainLock[31:0]].

IOMMUx4018 DeviceID Bank Lock Low

Bits	Description
31:0	DevIDLock[31:0]: deviceID lock enable bits[31:0] . Read-write. Reset: 0. DevIDLock[63:0] = {IOMMUx401C[DevIDLock[63:32]], DevIDLock[31:0]}. For each bit in DevIDLock[63:0], if the bit is set then writes to the corresponding bank in IOMMUx4[1,0][3:0]20 and IOMMUx4[1,0][3:0]24 are ignored. Bit positions above the value reported in IOMMUx4000[NCounterBanks] are ignored when written and return zero when read.

IOMMUx401C DeviceID Bank Lock High

Bits	Description
31:0	DevIDLock[63:32] . See: IOMMUx4018[DevIDLock[31:0]].

IOMMUx4[1,0][3:0]00 Counter Low**Table 184: Block to register mapping for IOMMUx4[1,0][3:0]00**

Register	Function	Register	Function
IOMMUx40000	Bank 0 Counter 0	IOMMUx41000	Bank 1 Counter 0
IOMMUx40100	Bank 0 Counter 1	IOMMUx41100	Bank 1 Counter 1
IOMMUx40200	Bank 0 Counter 2	IOMMUx41200	Bank 1 Counter 2
IOMMUx40300	Bank 0 Counter 3	IOMMUx41300	Bank 1 Counter 3

Bits	Description
31:0	Icounter[31:0] . Read-write. Reset: 0. Icounter[47:0] = {IOMMUx4[1,0][3:0]04[Icounter[47:32]], Icounter[31:0]}. Icounter[47:0] reports the counter value. The counter counts up continuously, wrapping at the maximum value. There is no overflow indicator.

IOMMUx4[1,0][3:0]04 Counter High**Table 185: Block to register mapping for IOMMUx4[1,0][3:0]04**

Register	Function	Register	Function
IOMMUx40004	Bank 0 Counter 0	IOMMUx41004	Bank 1 Counter 0
IOMMUx40104	Bank 0 Counter 1	IOMMUx41104	Bank 1 Counter 1
IOMMUx40204	Bank 0 Counter 2	IOMMUx41204	Bank 1 Counter 2
IOMMUx40304	Bank 0 Counter 3	IOMMUx41304	Bank 1 Counter 3

Bits	Description
31:16	Reserved.
15:0	Icounter [47:32]. See: IOMMUx4[1,0][3:0]00[Icounter[31:0]].

IOMMUx4[1,0][3:0]08 Counter Source**Table 186: Block to register mapping for IOMMUx4[1,0][3:0]08**

Register	Function	Register	Function
IOMMUx40008	Bank 0 Counter 0	IOMMUx41008	Bank 1 Counter 0
IOMMUx40108	Bank 0 Counter 1	IOMMUx41108	Bank 1 Counter 1
IOMMUx40208	Bank 0 Counter 2	IOMMUx41208	Bank 1 Counter 2
IOMMUx40308	Bank 0 Counter 3	IOMMUx41308	Bank 1 Counter 3

Bits	Description
31	Cac: counter source architectural or custom. Read-write. Reset: 0. 0=Architectural counter input group. 1=Custom input group.
30	CountUnits. Read-write. Reset: 0. 0=Counter counts events (level). 1=Counter counts clocks (edges).
29:8	Reserved.
7:0	Csource: counter source. Read-write. Reset: 0. Counter source. Selects event counter input from the choices provided.

IOMMUx4[1,0][3:0]10 PASID Match Low

See IOMMUx4008.

Table 187: Block to register mapping for IOMMUx4[1,0][3:0]10

Register	Function	Register	Function
IOMMUx40010	Bank 0 Counter 0	IOMMUx41010	Bank 1 Counter 0
IOMMUx40110	Bank 0 Counter 1	IOMMUx41110	Bank 1 Counter 1
IOMMUx40210	Bank 0 Counter 2	IOMMUx41210	Bank 1 Counter 2
IOMMUx40310	Bank 0 Counter 3	IOMMUx41310	Bank 1 Counter 3

Bits	Description
31	PasMEN: PASID match enable. Read-write. Reset: 0. 0=PASID is ignored. 1=Filtered PASID must match to count an event. An event with no PASID tag is only counted when Pasmen=0.
30:16	Reserved.
15:0	PasidMatch. Read-write. Reset: 0. This value is compared with the masked (filtered) value of the incoming PASID of the transaction to decide if the corresponding event is counted. The event is counted if PasidMatch is exactly equal to the masked incoming PASID.

IOMMUx4[1,0][3:0]14 PASID Match High

See [IOMMUx4008](#).

Table 188: Block to register mapping for IOMMUx4[1,0][3:0]14

Register	Function	Register	Function
IOMMUx40014	Bank 0 Counter 0	IOMMUx41014	Bank 1 Counter 0
IOMMUx40114	Bank 0 Counter 1	IOMMUx41114	Bank 1 Counter 1
IOMMUx40214	Bank 0 Counter 2	IOMMUx41214	Bank 1 Counter 2
IOMMUx40314	Bank 0 Counter 3	IOMMUx41314	Bank 1 Counter 3

Bits	Description
31:16	Reserved.
15:0	PasidMask. Read-write. Reset: 0. This bit-mask is ANDed with the PASID of the transaction to decide to count the corresponding event. 0=Count events for all values of incoming PASID. 0001h-FFFFh=Bit-wise mask ANDed with incoming PASID.

IOMMUx4[1,0][3:0]18 Domain Match Low

Table 189: Block to register mapping for IOMMUx4[1,0][3:0]18

Register	Function	Register	Function
IOMMUx40018	Bank 0 Counter 0	IOMMUx41018	Bank 1 Counter 0
IOMMUx40118	Bank 0 Counter 1	IOMMUx41118	Bank 1 Counter 1
IOMMUx40218	Bank 0 Counter 2	IOMMUx41218	Bank 1 Counter 2
IOMMUx40318	Bank 0 Counter 3	IOMMUx41318	Bank 1 Counter 3

Bits	Description
31	DomMEN: domain match enable. Read-write. Reset: 0. 0=Domain is ignored. 1=Filtered Domain must match DomainMatch to count an event.
30:16	Reserved.
15:0	DomainMatch. Read-write. Reset: 0. This value is compared with the masked (filtered) value of the incoming Domain of the transaction to decide to count the corresponding event. The event is counted if DomainMatch is exactly equal to the masked incoming Domain.

IOMMUx4[1,0][3:0]1C Domain Match High**Table 190: Block to register mapping for IOMMUx4[1,0][3:0]1C**

Register	Function	Register	Function
IOMMUx4001C	Bank 0 Counter 0	IOMMUx4101C	Bank 1 Counter 0
IOMMUx4011C	Bank 0 Counter 1	IOMMUx4111C	Bank 1 Counter 1
IOMMUx4021C	Bank 0 Counter 2	IOMMUx4121C	Bank 1 Counter 2
IOMMUx4031C	Bank 0 Counter 3	IOMMUx4131C	Bank 1 Counter 3

Bits	Description
31:16	Reserved.
15:0	DomainMask. Read-write. Reset: 0. This bit-mask is ANDed with the Domain of the transaction to decide to count the corresponding event. 0000h=Count events for all values of incoming Domain. 0001h-FFFFh=Bit-wise mask ANDed with incoming Domain.

IOMMUx4[1,0][3:0]20 DeviceID Match Low**Table 191: Block to register mapping for IOMMUx4[1,0][3:0]20**

Register	Function	Register	Function
IOMMUx40020	Bank 0 Counter 0	IOMMUx41020	Bank 1 Counter 0
IOMMUx40120	Bank 0 Counter 1	IOMMUx41120	Bank 1 Counter 1
IOMMUx40220	Bank 0 Counter 2	IOMMUx41220	Bank 1 Counter 2
IOMMUx40320	Bank 0 Counter 3	IOMMUx41320	Bank 1 Counter 3

Bits	Description
31	DidMEen: deviceID match enable. Read-write. Reset: 0. 0=DeviceID is ignored. 1=Filtered DeviceID must match to count an event.
30:16	Reserved.
15:0	DeviceidMatch. Read-write. Reset: 0. This value is compared with the masked (filtered) value of the incoming DeviceID of the transaction to decide to count the corresponding event. The event is counted if DeviceidMatch is exactly equal to the masked incoming DeviceID.

IOMMUx4[1,0][3:0]24 DeviceID Match High**Table 192: Block to register mapping for IOMMUx4[1,0][3:0]24**

Register	Function	Register	Function
IOMMUx40024	Bank 0 Counter 0	IOMMUx41024	Bank 1 Counter 0
IOMMUx40124	Bank 0 Counter 1	IOMMUx41124	Bank 1 Counter 1
IOMMUx40224	Bank 0 Counter 2	IOMMUx41224	Bank 1 Counter 2
IOMMUx40324	Bank 0 Counter 3	IOMMUx41324	Bank 1 Counter 3

Bits	Description
31:16	Reserved.
15:0	DeviceidMask . Read-write. Reset: 0. This bit-mask is ANDed with the DeviceID of the transaction to decide to count the corresponding event. 0=Count events for all values of incoming DeviceID. 0001h-FFFFh= Bit-wise mask ANDed with incoming DeviceID.

IOMMUx4[1,0][3:0]28 Counter Report Low

Table 193: Block to register mapping for IOMMUx4[1,0][3:0]28

Register	Function	Register	Function
IOMMUx40028	Bank 0 Counter 0	IOMMUx41028	Bank 1 Counter 0
IOMMUx40128	Bank 0 Counter 1	IOMMUx41128	Bank 1 Counter 1
IOMMUx40228	Bank 0 Counter 2	IOMMUx41228	Bank 1 Counter 2
IOMMUx40328	Bank 0 Counter 3	IOMMUx41328	Bank 1 Counter 3

Bits	Description
31:0	EventNote[31:0] . Read-write. Reset: 0. EventNote[51:0] = {IOMMUx4[1,0][3:0]2C[EventNode[51:32]], EventNode[31:0]}. When IOMMUx4[1,0][3:0]2C[CERE]=1 and the corresponding counter is incremented and wraps to zero, EventNote[51:0] is reported in the EVENT_COUNTER_ZERO event log entry.

IOMMUx4[1,0][3:0]2C Counter Report High

Table 194: Block to register mapping for IOMMUx4[1,0][3:0]2C

Register	Function	Register	Function
IOMMUx4002C	Bank 0 Counter 0	IOMMUx4102C	Bank 1 Counter 0
IOMMUx4012C	Bank 0 Counter 1	IOMMUx4112C	Bank 1 Counter 1
IOMMUx4022C	Bank 0 Counter 2	IOMMUx4122C	Bank 1 Counter 2
IOMMUx4032C	Bank 0 Counter 3	IOMMUx4132C	Bank 1 Counter 3

Bits	Description
31	CERE: counter event report enable . Read-write. Reset: 0. Counter Event Report Enable. 0=no event report when counter wraps to zero. 1=IOMMU writes an EVENT_COUNTER_ZERO event log entry when the counter wraps to zero. The counter-wrap event is treated like any other event.. Software note: the counter-wrap event is delivered promptly but without a latency guarantee.
30:20	Reserved.
19:0	EventNote[51:32] . See: IOMMUx4[1,0][3:0]28[EventNote[31:0]].

3.16 APIC Registers

See 2.4.8.1.2 [APIC Register Space].

MMIO local APIC space is accessible in xAPIC mode.

APIC20 APIC ID

Bits	Description
31:24	ApicId: APIC ID. Read-write. Reset: Varies based on core number. The initial value of APIC20[ApicId[7:0]] is {0b, D18F0x60[NodeId[2:0]], CpuCoreNum[3:0]}. See 2.4.8.1.3 [ApicId Enumeration Requirements]. See 2.4.3 [Processor Cores and Downcoring].
23:0	Reserved.

APIC30 APIC Version

Read-only.

Bits	Description
31	ExtApicSpace: extended APIC register space present. Reset: 1. 1=Indicates the presence of extended APIC register space starting at APIC400.
30:25	RAZ.
24	DirectedEoiSupport: directed EOI support. Reset: 0. 0=Directed EOI capability not supported.
23:16	MaxLvtEntry. Reset: Product-specific. Specifies the number of entries in the local vector table minus one.
15:8	RAZ.
7:0	Version. Reset: 10h. Indicates the version number of this APIC implementation.

APIC80 Task Priority (TPR)

Bits	Description
31:8	RAZ.
7:0	Priority. Read-write. Reset: 0. This field is assigned by software to set a threshold priority at which the core is interrupted.

APIC90 Arbitration Priority (APR)

Bits	Description
31:8	RAZ.
7:0	Priority. Read-only. Reset: 0. Indicates the current priority for a pending interrupt, or a task or interrupt being serviced by the core. The priority is used to arbitrate between cores to determine which accepts a lowest-priority interrupt request.

APICA0 Processor Priority (PPR)

Bits	Description
31:8	RAZ.
7:0	Priority. Read-only. Reset: 0. Indicates the core's current priority servicing a task or interrupt, and is used to determine if any pending interrupts should be serviced. It is the higher value of the task priority value and the current highest in-service interrupt.

APICB0 End of Interrupt

This register is written by the software interrupt handler to indicate the servicing of the current interrupt is complete.

Bits	Description
31:0	Reserved. Write-only.

APICC0 Remote Read

Reset: 0.

Bits	Description
31:0	RemoteReadData. Read-only. The data resulting from a valid completion of a remote read inter-processor interrupt.

APICD0 Logical Destination (LDR)

Reset: 0.

Bits	Description
31:24	Destination. Read-write. This APIC's destination identification. Used to determine which interrupts should be accepted.
23:0	Reserved.

APICE0 Destination Format

Reset: FFFF_FFFFh.

Bits	Description								
31:28	Format. Read-write. Controls which format to use when accepting interrupts with a logical destination mode. <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><u>Bits</u></td> <td><u>Definition</u></td> </tr> <tr> <td>0h</td> <td>Cluster destinations are used</td> </tr> <tr> <td>Eh-1h</td> <td>Reserved</td> </tr> <tr> <td>Fh</td> <td>Flat destinations are used</td> </tr> </table>	<u>Bits</u>	<u>Definition</u>	0h	Cluster destinations are used	Eh-1h	Reserved	Fh	Flat destinations are used
<u>Bits</u>	<u>Definition</u>								
0h	Cluster destinations are used								
Eh-1h	Reserved								
Fh	Flat destinations are used								
27:0	Reserved.								

APICF0 Spurious-Interrupt Vector (SVR)

Bits	Description
31:13	RAZ.
12	EoiBroadcastDisable: EOI broadcast disable. Read-only. Reset: 0.
11:10	RAZ.
9	FocusDisable. Read-write. Reset: 0. 1=Disable focus core checking during lowest-priority arbitrated interrupts.
8	APICSWEn: APIC software enable. Read-write. Reset: 0. 0=SMI, NMI, INIT, LINT[1:0], and Startup interrupts may be accepted; pending interrupts in APIC[170:100] and APIC[270:200] are held, but further fixed, lowest-priority, and ExtInt interrupts are not accepted. All LVT entry mask bits are set and cannot be cleared.
7:0	Vector. Read-write. Reset: FFh. The vector that is sent to the core in the event of a spurious interrupt. The behavior of bits 3:0 are controlled as specified by D18F0x68 [Link Transaction Control][Api-cExtSpur] .

APIC[170:100] In-Service (ISR)

Reset: 0. Read-only.

The in-service registers provide a bit per interrupt to indicate that the corresponding interrupt is being serviced by the core. APIC100[15:0] are reserved. Interrupts are mapped as follows:

Table 195: Register Mapping for APIC[170:100]

Register	Function
APIC100	Interrupts [31:16]
APIC110	Interrupts [63:32]
APIC120	Interrupts [95:64]
APIC130	Interrupts [127:96]
APIC140	Interrupts [159:128]
APIC150	Interrupts [191:160]
APIC160	Interrupts [223:192]
APIC170	Interrupts [255:224]

Bits	Description
31:0	InServiceBits. These bits are set when the corresponding interrupt is being serviced by the core.

APIC[1F0:180] Trigger Mode (TMR)

Reset: 0. Read-only. The trigger mode registers provide a bit per interrupt to indicate the assertion mode of each interrupt. APIC180[15:0] are reserved. Interrupts are mapped as follows:

Table 196: Register Mapping for APIC[1F0:180]

Register	Function
APIC180	Interrupts [31:16]
APIC190	Interrupts [63:32]
APIC1A0	Interrupts [95:64]
APIC1B0	Interrupts [127:96]
APIC1C0	Interrupts [159:128]
APIC1D0	Interrupts [191:160]
APIC1E0	Interrupts [223:192]
APIC1F0	Interrupts [255:224]

Bits	Description
31:0	TriggerModeBits. The corresponding trigger mode bit is updated when an interrupt is accepted. The values are: 0=Edge-triggered interrupt. 1=Level-triggered interrupt.

APIC[270:200] Interrupt Request (IRR)

Reset: 0. Read-only.

The interrupt request registers provide a bit per interrupt to indicate that the corresponding interrupt has been accepted by the APIC. APIC200[15:0] are reserved. Interrupts are mapped as follows:

Table 197: Register Mapping for APIC[270:200]

Register	Function
APIC200	Interrupts [31:16]
APIC210	Interrupts [63:32]
APIC220	Interrupts [95:64]
APIC230	Interrupts [127:96]
APIC240	Interrupts [159:128]
APIC250	Interrupts [191:160]
APIC260	Interrupts [223:192]
APIC270	Interrupts [255:224]

Bits	Description
31:0	RequestBits. The corresponding request bit is set when the an interrupt is accepted by the APIC.

APIC280 Error Status

Writes to this register trigger an update of the register state. The value written by software is arbitrary. Each write causes the internal error state to be loaded into this register, clearing the internal error state. Consequently, a second write prior to the occurrence of another error causes the register to be overwritten with cleared data.

Bits	Description
31:8	RAZ.

7	IllegalRegAddr: illegal register address. Read-write. Reset: 0. This bit indicates that an access to a nonexistent register location within this APIC was attempted. Can only be set in xAPIC mode.
6	RcvdIllegalVector: received illegal vector. Read-write. Reset: 0. This bit indicates that this APIC has received a message with an illegal vector (00h to 0Fh for fixed and lowest priority interrupts).
5	SentIllegalVector. Read-write. Reset: 0. This bit indicates that this APIC attempted to send a message with an illegal vector (00h to 0Fh for fixed and lowest priority interrupts).
4	RAZ.
3	RcvAcceptError: receive accept error. Read-write. Reset: 0. This bit indicates that a message received by this APIC was not accepted by this or any other APIC.
2	SendAcceptError. Read-write. Reset: 0. This bit indicates that a message sent by this APIC was not accepted by any APIC.
1:0	RAZ.

APIC300 Interrupt Command Low (ICR Low)

Not all combinations of ICR fields are valid. Only the following combinations are valid:

Table 198: ICR valid combinations

Message Type	Trigger Mode	Level	Destination Shorthand
Fixed	Edge	x	x
	Level	Assert	x
Lowest Priority, SMI, NMI, INIT	Edge	x	Destination or all excluding self.
	Level	Assert	Destination or all excluding self
Startup	x	x	Destination or all excluding self

Note: x indicates a don't care.

Bits	Description										
31:20	RAZ.										
19:18	<p>DestShrthnd: destination shorthand. Read-write. Reset: 0. Provides a quick way to specify a destination for a message.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No shorthand (Destination field)</td> </tr> <tr> <td>01b</td> <td>Self</td> </tr> <tr> <td>10b</td> <td>All including self</td> </tr> <tr> <td>11b</td> <td>All excluding self (This sends a message with a destination encoding of all 1s, so if lowest priority is used the message could end up being reflected back to this APIC.)</td> </tr> </tbody> </table> <p>If all including self or all excluding self is used, then destination mode is ignored and physical is automatically used.</p>	<u>Bits</u>	<u>Description</u>	00b	No shorthand (Destination field)	01b	Self	10b	All including self	11b	All excluding self (This sends a message with a destination encoding of all 1s, so if lowest priority is used the message could end up being reflected back to this APIC.)
<u>Bits</u>	<u>Description</u>										
00b	No shorthand (Destination field)										
01b	Self										
10b	All including self										
11b	All excluding self (This sends a message with a destination encoding of all 1s, so if lowest priority is used the message could end up being reflected back to this APIC.)										

17:16	RemoteRdStat: remote read status. Read-only. Reset: 0. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Read was invalid</td> </tr> <tr> <td>01b</td> <td>Delivery pending</td> </tr> <tr> <td>10b</td> <td>Delivery complete and access was valid</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	Read was invalid	01b	Delivery pending	10b	Delivery complete and access was valid	11b	Reserved								
Bits	Description																		
00b	Read was invalid																		
01b	Delivery pending																		
10b	Delivery complete and access was valid																		
11b	Reserved																		
15	TM: trigger mode. Read-write. Reset: 0. Indicates how this interrupt is triggered. 0=Edge triggered. 1=Level triggered.																		
14	Level. Read-write. Reset: 0. 0=Deasserted. 1=Asserted.																		
13	RAZ.																		
12	DS: interrupt delivery status. Read-only. Reset: 0. In xAPIC mode this bit is set to indicate that the interrupt has not yet been accepted by the destination core(s). 0=Idle. 1=Send pending. Software may repeatedly write ICRL without polling the DS bit; all requested IPIs will be delivered.																		
11	DM: destination mode. Read-write. Reset: 0. 0=Physical. 1=Logical.																		
10:8	MsgType. Read-write. Reset: 0. The message types are encoded as follows: <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Fixed</td> </tr> <tr> <td>001b</td> <td>Lowest Priority.</td> </tr> <tr> <td>010b</td> <td>SMI</td> </tr> <tr> <td>011b</td> <td>Remote read.</td> </tr> <tr> <td>100b</td> <td>NMI</td> </tr> <tr> <td>101b</td> <td>INIT</td> </tr> <tr> <td>110b</td> <td>Startup</td> </tr> <tr> <td>111b</td> <td>External interrupt.</td> </tr> </tbody> </table>	Bits	Description	000b	Fixed	001b	Lowest Priority.	010b	SMI	011b	Remote read.	100b	NMI	101b	INIT	110b	Startup	111b	External interrupt.
Bits	Description																		
000b	Fixed																		
001b	Lowest Priority.																		
010b	SMI																		
011b	Remote read.																		
100b	NMI																		
101b	INIT																		
110b	Startup																		
111b	External interrupt.																		
7:0	Vector. Read-write. Reset: 0. The vector that is sent for this interrupt source.																		

APIC310 Interrupt Command High (ICR High)

Bits	Description
31:24	DestinationField. Read-write. Reset: 0. The destination encoding used when APIC300[DestShrthnd] is 00b.
23:0	RAZ.

APIC320 LVT Timer

Bits	Description
31:18	RAZ.
17	Mode. Read-write. Reset: 0. 0=One-shot. 1=Periodic.
16	Mask. Read-write. Reset: 1. 0=Not masked. 1=Masked.
15:13	RAZ.
12	DS: interrupt delivery status. Read-only; updated-by-hardware. Reset: 0. 0=Idle. 1=Send pending. (Indicates that the interrupt has not yet been accepted by the core.)
11	RAZ.

10:8	MsgType: message type. Read-write. Reset: 000b. See 2.4.8.1.14 [Generalized Local Vector Table].
7:0	Vector. Read-write. Reset: 00h. Interrupt vector number.

APIC330 LVT Thermal Sensor

Interrupts for this local vector table are caused by changes in [MSRC001_0061 \[P-state Current Limit\]\[CurP-stateLimit\]](#) due to SB-RMI or HTC.

Bits	Description
31:17	RAZ.
16	Mask. Read-write. Reset: 1. 0=Not masked. 1=Masked.
15:13	RAZ.
12	DS: interrupt delivery status. Read-only; updated-by-hardware. Reset: 0. 0=Idle. 1=Send pending. (Indicates that the interrupt has not yet been accepted by the core.)
11	RAZ.
10:8	MsgType: message type. Read-write. Reset: 000b. See 2.4.8.1.14 [Generalized Local Vector Table].
7:0	Vector. Read-write. Reset: 00h. Interrupt vector number.

APIC340 LVT Performance Monitor

Interrupts for this local vector table are caused by overflows of

- [MSRC001_00\[07:04\] \[Performance Event Counter \(PERF_CTR\[3:0\]\)\]](#).
- [MSRC001_020\[B,9,7,5,3,1\] \[Performance Event Counter \(PERF_CTR\[5:0\]\)\]](#).
- [MSRC001_024\[7,5,3,1\] \[Northbridge Performance Event Counter \(NB_PERF_CTR\[3:0\]\)\]](#).

The Mask bit is not set automatically when the interrupt is taken.

Bits	Description
31:17	RAZ.
16	Mask. Read-write. Reset: 1. 0=Not masked. 1=Masked.
15:13	RAZ.
12	DS: interrupt delivery status. Read-only; updated-by-hardware. Reset: 0. 0=Idle. 1=Send pending. (Indicates that the interrupt has not yet been accepted by the core.)
11	RAZ.
10:8	MsgType: message type. Read-write. Reset: 000b. See 2.4.8.1.14 [Generalized Local Vector Table].
7:0	Vector. Read-write. Reset: 00h. Interrupt vector number.

APIC3[60:50] LVT LINT[1:0]

Table 199: Register Mapping for APIC3[60:50]

Register	Function
APIC350	LINT 0
APIC360	LINT 1

Bits	Description
31:17	RAZ.
16	Mask. Read-write. Reset: 1. 0=Not masked. 1=Masked.
15	TM: trigger mode. Read-write. Reset: 0. 0=Edge. 1=Level.
14	RmtIRR. Read-only; updated-by-hardware. Reset: 0. If trigger mode is level, remote IRR is set when the interrupt has begun service. Remote IRR is cleared when the end of interrupt has occurred.
13	Reserved.
12	DS: interrupt delivery status. Read-only; updated-by-hardware. Reset: 0. 0=Idle. 1=Send pending. (Indicates that the interrupt has not yet been accepted by the core.)
11	RAZ.
10:8	MsgType: message type. Read-write. Reset: 000b. See 2.4.8.1.14 [Generalized Local Vector Table].
7:0	Vector. Read-write. Reset: 00h. Interrupt vector number.

APIC370 LVT Error

Bits	Description
31:17	RAZ.
16	Mask. Read-write. Reset: 1. 0=Not masked. 1=Masked.
15:13	Reserved.
12	DS: interrupt delivery status. Read-only; updated-by-hardware. Reset: 0. 0=Idle. 1=Send pending. (Indicates that the interrupt has not yet been accepted by the core.)
11	RAZ.
10:8	MsgType: message type. Read-write. Reset: 000b. See 2.4.8.1.14 [Generalized Local Vector Table].
7:0	Vector. Read-write. Reset: 00h. Interrupt vector number.

APIC380 Timer Initial Count

Reset: 0.

Bits	Description
31:0	Count. Read-write. The value copied into the current count register when the timer is loaded or reloaded.

APIC390 Timer Current Count

Reset: 0.

Bits	Description
31:0	Count. Read-only. The current value of the counter.

APIC3E0 Timer Divide Configuration

The Div bits are encoded as follows:

Table 200: Div[3,1:0] Value Table

Div[3]	Div[1:0]	Resulting Timer Divide
0	00b	2
0	01b	4
0	10b	8
0	11b	16
1	00b	32
1	01b	64
1	10b	128
1	11b	1

Bits	Description
31:4	RAZ.
3	Div[3] . Read-write. Reset: 0. See Table 200 .
2	RAZ.
1:0	Div[1:0] . Read-write. Reset: 0. See Table 200 .

APIC400 Extended APIC Feature

Bits	Description
31:24	RAZ.
23:16	ExtLvtCount: extended local vector table count . Read-only. Reset: 04h. This specifies the number of extended LVT registers (APIC[530:500]) in the local APIC.
15:3	RAZ.
2	ExtApicIdCap: extended APIC ID capable . Read-only. Reset: 1. 1=The processor is capable of supporting an 8-bit APIC ID, as controlled by APIC410[ExtApicIdEn] .
1	SeoiCap: specific end of interrupt capable . Read-only. Reset: 1. 1=The APIC420 [Specific End Of Interrupt] is present.
0	IerCap: interrupt enable register capable . Read-only. Reset: 1. This bit indicates that the APIC[4F0:480] [Interrupt Enable] are present. See 2.4.8.1.8 [Interrupt Masking] .

APIC410 Extended APIC Control

Bits	Description
31:3	RAZ.
2	ExtApicIdEn: extended APIC ID enable . Read-write. Reset: 0. 1=Enable 8-bit APIC ID; APIC20[ApicId] supports an 8-bit value; an interrupt broadcast in physical destination mode requires that the $\text{IntDest}[7:0]=1111_1111\text{b}$ (instead of $\text{xxxx_1111}\text{b}$); a match in physical destination mode occurs when $(\text{IntDest}[7:0] == \text{ApicId}[7:0])$ instead of $(\text{IntDest}[3:0] == \text{ApicId}[3:0])$. If $\text{ExtApicIdEn}=1$ then program D18F0x68[ApicExtId]=1 and D18F0x68[ApicExtBrdCst]=1 .

1	SeoiEn. Read-write. Reset: 0. 1=Enable SEOI generation when a write to APIC420 [Specific End Of Interrupt] is received.
0	IerEn. Read-write. Reset: 0. 1=Enable writes to the interrupt enable registers.

APIC420 Specific End Of Interrupt

Bits	Description
31:8	RAZ.
7:0	EoiVec: end of interrupt vector. Read-write. Reset: 0. A write to this field causes an end of interrupt cycle to be performed for the vector specified in this field. The behavior is undefined if no interrupt is pending for the specified interrupt vector.

APIC[4F0:480] Interrupt Enable

Reset: FFFF_FFFFh. Interrupt enables range is mapped as follows:

Table 201: Register Mapping for APIC[4F0:480]

Register	Function
APIC480	IntEn[31:0]
APIC490	IntEn[63:32]
APIC4A0	IntEn[95:64]
APIC4B0	IntEn[127:96]
APIC4C0	IntEn[159:128]
APIC4D0	IntEn[191:160]
APIC4E0	IntEn[223:192]
APIC4F0	IntEn[255:224]

Bits	Description
31:0	InterruptEnableBits. Read-write. The interrupt enable bits can be used to enable each of the 256 interrupts. See above table.

APIC[530:500] Extended Interrupt [3:0] Local Vector Table

APIC500 provides a local vector table entry for IBS; See D18F3x1CC. The APIC[530:510] registers can be used by system BIOS to provide local vector table entries for internal interrupt sources.

Table 202: Register Mapping for APIC[530:500]

Register	Function
APIC500	Extended Interrupt 0 (IBS)
APIC510	Extended Interrupt 1 (Thresholding)
APIC520	Extended Interrupt 2 (Unused)
APIC530	Extended Interrupt 3 (SBI)

Bits	Description
31:17	RAZ.
16	Mask. Read-write. Reset: 1. 0=Not masked. 1=Masked.
15:13	MBZ.
12	DS: interrupt delivery status. Read-only; updated-by-hardware. Reset: 0. 0=Idle. 1=Send pending. (Indicates that the interrupt has not yet been accepted by the core.)
11	RAZ.
10:8	MsgType: message type. Read-write. Reset: 000b. See 2.4.8.1.14 [Generalized Local Vector Table] .
7:0	Vector. Read-write. Reset: 00h. Interrupt vector number.

3.17 CPUID Instruction Registers

Processor feature capabilities and configuration information are provided through the CPUID instruction. The information is accessed by (1) selecting the CPUID function setting EAX and optionally ECX for some functions, (2) executing the CPUID instruction, and (3) reading the results in the EAX, EBX, ECX, and EDX registers. The syntax *CPUID FnXXXX_XXXX_EiX[_xYYY]* refers to the function where EAX==X, and optionally ECX==Y, and the registers specified by EiX. EiX can be any single register such as {EAX, EBX, ECX, and EDX}, or a range of registers, such as E[C,B,A]X. Undefined function numbers return 0's in all 4 registers. See [2.4.10 \[CPUID Instruction\]](#).

Unless otherwise specified, single-bit feature fields are encoded as 1=Feature is supported by the processor; 0=Feature is not supported by the processor.

The following provides AMD Family 15h Models 10h-1Fh processor specific details about CPUID. See the *CPUID Specification* for further information.

CPUID Fn0000_0000_EAX Processor Vendor and Largest Standard Function Number

Bits	Description
31:0	LFuncStd: largest standard function. Value: 0000_000Dh. The largest CPUID standard function input value supported by the processor implementation.

CPUID Fn0000_0000_E[D,C,B]X Processor Vendor

[CPUID Fn0000_0000_E\[D,C,B\]X](#) and [CPUID Fn8000_0000_E\[D,C,B\]X](#) return the same value.

Table 203: CPUID Fn0000_0000_E[B,C,D]X Value

Register	Value	Description
CPUID Fn0000_0000_EBX	6874_7541h	The ASCII characters “h t u A”.
CPUID Fn0000_0000_ECX	444D_4163h	The ASCII characters “D M A c”.
CPUID Fn0000_0000_EDX	6974_6E65h	The ASCII characters “i t n e”.

Bits	Description
31:0	Vendor. The 12 8-bit ASCII character codes to create the string “AuthenticAMD”.

CPUID Fn0000_0001_EAX Family, Model, Stepping Identifiers

[CPUID Fn0000_0001_EAX](#), [CPUID Fn8000_0001_EAX](#) are an alias of [D18F3xFC](#).

Family is an 8-bit value and is defined as: **Family[7:0]** = ({0000b,BaseFamily[3:0]} + ExtendedFamily[7:0]). E.g. If BaseFamily[3:0]=Fh and ExtendedFamily[7:0]=07h, then Family[7:0]=16h.

Model is an 8-bit value and is defined as: **Model[7:0]** = {ExtendedModel[3:0], BaseModel[3:0]}. E.g. If ExtendedModel[3:0]=Eh and BaseModel[3:0]=8h, then Model[7:0] = E8h. Model numbers vary with product.

Bits	Description
31:0	Alias of D18F3xFC .

CPUID Fn0000_0001_EBX LocalApicId, LogicalProcessorCount, CLFlush

Bits	Description
31:24	LocalApicId: initial local APIC physical ID. The initial APIC20[ApicId] value. After D18F0x60[NodeId] has been initialized, changes to APIC20[ApicId] do not affect the value of this CPUID register. See 2.4.3 [Processor Cores and Downcoring].
23:16	LogicalProcessorCount: logical processor count. If CPUID Fn0000_0001_EDX[HTT] = 1, then this field indicates the number of cores in the processor. Value: IF (CPUID Fn0000_0001_EDX[HTT]) THEN CPUID Fn8000_0008_ECX[NC] + 1 ELSE 0 ENDIF.
15:8	CLFlush: CLFLUSH size in quadwords. Value: 08h.
7:0	8BitBrandId: 8 bit brand ID. Value: 00h. Indicates that the brand ID is in CPUID Fn8000_0001_EBX.

CPUID Fn0000_0001_ECX Feature Identifiers

These values can be over-written by MSRC001_1004.

Bits	Description
31	RAZ. Reserved for use by hypervisor to indicate guest status.
30	Reserved.
29	F16C: half-precision convert instruction support. Value: 1.
28	AVX: AVX instruction support. Value: 1.
27	OSXSAVE: OS enabled support for XGETBV/XSETBV. 1=The OS has enabled support for XGETBV/XSETBV instructions to query processor extended states.
26	XSAVE: XSAVE (and related) instruction support. Value: 1. 1=Support provided for the XSAVE, XRSTOR, XSETBV, and XGETBV instructions and the XFEATURE_ENABLED_MASK register.
25	AES: AES instruction support. Value: Product-specific.
24	Reserved.
23	POPCNT: POPCNT instruction. Value: 1.
22	Reserved.
21	x2APIC: x2APIC capability. Value: 0.
20	SSE42: SSE4.2 instruction support. Value: 1.
19	SSE41: SSE4.1 instruction support. Value: 1.
18:14	Reserved.
13	CMPXCHG16B: CMPXCHG16B instruction. Value: 1.
12	FMA: FMA instruction support. Value: 1.
11:10	Reserved.
9	SSSE3: supplemental SSE3 extensions. Value: 1.
8:4	Reserved.
3	Monitor: Monitor/Mwait instructions. Value: ~MSRC001_0015[MonMwaitDis].
2	Reserved.

Bits	Description
1	PCLMULQDQ: PCLMULQDQ instruction support. Value: Product-specific.
0	SSE3: SSE3 extensions. Value: 1.

CPUID Fn0000_0001_EDX Feature Identifiers

These values can be over-written by [MSRC001_1004](#).

Bits	Description
31:29	Reserved.
28	HTT: hyper-threading technology. Value: CPUID Fn8000_0008_ECX[NC] !=0. This bit qualifies the meaning of CPUID Fn0000_0001_EBX[LogicalProcessorCount] . 1=Multi core product (CPUID Fn8000_0008_ECX[NC] != 0). 0=Single core product (CPUID Fn8000_0008_ECX[NC] ==0).
27	Reserved.
26	SSE2: SSE2 extensions. Value: 1.
25	SSE: SSE extensions. Value: 1.
24	FXSR: FXSAVE and FXRSTOR instructions. Value: 1.
23	MMX: MMX™ instructions. Value: 1.
22:20	Reserved.
19	CLFSH: CLFLUSH instruction. Value: 1.
18	Reserved.
17	PSE36: page-size extensions. Value: 1.
16	PAT: page attribute table. Value: 1.
15	CMOV: conditional move instructions, CMOV, FCOMI, FCMOV. Value: 1.
14	MCA: machine check architecture, MCG_CAP. Value: 1.
13	PGE: page global extension, CR4.PGE. Value: 1.
12	MTRR: memory-type range registers. Value: 1.
11	SysEnterSysExit: SYSENTER and SYSEXIT instructions. Value: 1.
10	Reserved.
9	APIC: advanced programmable interrupt controller (APIC) exists and is enabled. Value: MSR0000_001B[ApicEn] .
8	CMPXCHG8B: CMPXCHG8B instruction. Value: 1.
7	MCE: machine check exception, CR4.MCE. Value: 1.
6	PAE: physical-address extensions (PAE). Value: 1.
5	MSR: AMD model-specific registers (MSRs), with RDMSR and WRMSR instructions. Value: 1.
4	TSC: time stamp counter, RDTSC/RDTSCP instructions, CR4.TSD. Value: 1.
3	PSE: page-size extensions (4 MB pages). Value: 1.
2	DE: debugging extensions, IO breakpoints, CR4.DE. Value: 1.
1	VME: virtual-mode enhancements. Value: 1.
0	FPU: x87 floating point unit on-chip. Value: 1.

CPUID Fn0000_000[4,3,2] Reserved

Bits	Description
31:0	Reserved.

CPUID Fn0000_0005_EAX Monitor/MWait

Bits	Description
31:16	Reserved.
15:0	MonLineSizeMin: smallest monitor-line size in bytes. Value: 40h.

CPUID Fn0000_0005_EBX Monitor/MWait

Bits	Description
31:16	Reserved.
15:0	MonLineSizeMax: largest monitor-line size in bytes. Value: 40h.

CPUID Fn0000_0005_ECX Monitor/MWait

Bits	Description
31:2	Reserved.
1	IBE: interrupt break-event. Value: 1.
0	EMX: enumerate MONITOR/MWAIT extensions. Value: 1.

CPUID Fn0000_0005_EDX Monitor/MWait

Bits	Description
31:0	Reserved.

CPUID Fn0000_0006_EAX Thermal and Power Management

Bits	Description
31:0	Reserved.

CPUID Fn0000_0006_EBX Thermal and Power Management

Bits	Description
31:0	Reserved.

CPUID Fn0000_0006_ECX Thermal and Power Management

These values can be over-written by [MSRC001_1003](#).

Bits	Description
31:1	Reserved.
0	EffFreq: effective frequency interface. Value: 1. 1=Indicates presence of MSR0000_00E7 [Max Performance Frequency Clock Count (MPERF)] and MSR0000_00E8 [Actual Performance Frequency Clock Count (APERF)] .

CPUID Fn0000_0006_EDX Thermal and Power Management

Bits	Description
31:0	Reserved.

CPUID Fn0000_0007_EAX_x0 Structured Extended Feature Identifiers (ECX=0)

Bits	Description
31:0	Reserved.

CPUID Fn0000_0007_EBX_x0 Structured Extended Feature Identifiers (ECX=0)

These values can be over-written by [MSRC001_1002](#).

Bits	Description
31:4	Reserved.
3	BMI1: bit manipulation instruction support. Value: 1.
2:0	Reserved.

CPUID Fn0000_0007_ECX_x0 Structured Extended Feature Identifiers (ECX=0)

Bits	Description
31:0	Reserved.

CPUID Fn0000_0007_EDX_x0 Structured Extended Feature Identifiers (ECX=0)

Bits	Description
31:0	Reserved.

CPUID Fn0000_000[A:8] Reserved

Bits	Description
31:0	Reserved.

CPUID Fn0000_000B_EAX Reserved

Bits	Description
31:0	Reserved.

CPUID Fn0000_000B_EBX Extended Topology Enumeration

Bits	Description
31:0	Reserved. Value: 0. 0=Extended topology enumeration leaf is not available.

CPUID Fn0000_000B_ECX Reserved

Bits	Description
31:0	Reserved.

CPUID Fn0000_000B_EDX Reserved

Bits	Description
31:0	Reserved.

CPUID Fn0000_000C Reserved

Bits	Description
31:0	Reserved.

CPUID Fn0000_000D_EAX_x0 Processor Extended State Enumeration (ECX=0)

Bits	Description
31:0	XFeatureSupportedMask[31:0] . Value: 0000_0007h.

CPUID Fn0000_000D_EBX_x0 Processor Extended State Enumeration (ECX=0)

Bits	Description
31:0	<p>XFeatureEnabledSizeMax. Size in bytes of XSAVE/XRSTOR area for the currently enabled features in XCR0. Value: 512 + 64 + (IF (XCR0[AVX] XCR0[LWP]) THEN 256 ELSE 0 ENDIF) + (IF XCR0[LWP]) THEN 128 ELSE 0 ENDIF). The components of this sum are described as follows:</p> <ul style="list-style-type: none"> • 512: FPU/SSE save area (needed even if XCR0[SSE]=0) • 64: Header size (always needed). • Size of YMM area if YMM enabled OR if LWP enabled. • Size of LWP area if LWP enabled.

CPUID Fn0000_000D_ECX_x0 Processor Extended State Enumeration (ECX=0)

Bits	Description
31:0	XFeatureSupportedSizeMax. Value: 0000_03C0h. Size in bytes of XSAVE/XRSTOR area for all features that the core supports. See XFeatureEnabledSizeMax.

CPUID Fn0000_000D_EDX_x0 Processor Extended State Enumeration (ECX=0)

Bits	Description
31:0	XFeatureSupportedMask[63:32]. Value: 4000_0000h.

CPUID Fn0000_000D_EAX_x2 Processor Extended State Enumeration (ECX=2)

Bits	Description
31:0	YmmSaveStateSize: YMM save state byte size. Value: 0000_0100h.

CPUID Fn0000_000D_EBX_x2 Processor Extended State Enumeration (ECX=2)

Bits	Description
31:0	YmmSaveStateOffset: YMM save state byte offset. Value: 0000_0240h.

CPUID Fn0000_000D_ECX_x2 Processor Extended State Enumeration (ECX=2)

Bits	Description
31:0	Reserved.

CPUID Fn0000_000D_EDX_x2 Processor Extended State Enumeration (ECX=2)

Bits	Description
31:0	Reserved.

For CPUID Fn0000_000D, if ECX>2 and ECX<62 then EAX/EBX/ECX/EDX will return 0.

CPUID Fn0000_000D_EAX_x3E Processor Extended State Enumeration (ECX=62)

Bits	Description
31:0	LwpSaveStateSize: LWP save state byte size. Value: 0000_0080h.

CPUID Fn0000_000D_EBX_x3E Processor Extended State Enumeration (ECX=62)

Bits	Description
31:0	LwpSaveStateOffset : LWP save state byte offset. Value: 0000_0340h.

CPUID Fn0000_000D_ECX_x3E Processor Extended State Enumeration (ECX=62)

Bits	Description
31:0	Reserved.

CPUID Fn0000_000D_EDX_x3E Processor Extended State Enumeration (ECX=62)

Bits	Description
31:0	Reserved.

For CPUID Fn0000_000D, if ECX>62 then EAX/EBX/ECX/EDX will return 0.

CPUID Fn8000_0000_EAX Largest Extended Function Number

Bits	Description
31:0	LFuncExt : largest extended function. Value: 8000_001Eh. The largest CPUID extended function input value supported by the processor implementation.

CPUID Fn8000_0000_E[D,C,B]X Processor Vendor

CPUID Fn0000_0000_E[D,C,B]X and CPUID Fn8000_0000_E[D,C,B]X return the same value.

Table 204: CPUID Fn8000_0000_E[B,C,D]X Value

Register	Value	Description
CPUID Fn8000_0000_EBX	6874_7541h	The ASCII characters “h t u A”.
CPUID Fn8000_0000_ECX	444D_4163h	The ASCII characters “D M A c”.
CPUID Fn8000_0000_EDX	6974_6E65h	The ASCII characters “i t n e”.

Bits	Description
31:0	Vendor . The 12 8-bit ASCII character codes to create the string “AuthenticAMD”.

CPUID Fn8000_0001_EAX Family, Model, Stepping Identifiers

See CPUID Fn0000_0001_EAX.

Bits	Description
31:0	Alias of D18F3xFC.

CPUID Fn8000_0001_EBX BrandId Identifier

Bits	Description								
31:28	PkgType: package type. Specifies the package type. Value: Product-specific. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>FP2 (BGA)</td> </tr> <tr> <td>0001b</td> <td>FS1r2 (uPGA)</td> </tr> <tr> <td>0010b</td> <td>FM2 (PGA)</td> </tr> </tbody> </table>	Bits	Description	0000b	FP2 (BGA)	0001b	FS1r2 (uPGA)	0010b	FM2 (PGA)
Bits	Description								
0000b	FP2 (BGA)								
0001b	FS1r2 (uPGA)								
0010b	FM2 (PGA)								
27:0	Reserved.								

CPUID Fn8000_0001_ECX Feature Identifiers

These values can be over-written by [MSRC001_1005](#).

Bits	Description
31:25	Reserved.
24	PerfCtrExtNB: NB performance counter extensions support. Value: 1. Indicates support for MSRC001_024[6,4,2,0] and MSRC001_024[7,5,3,1] .
23	PerfCtrExtCore: core performance counter extensions support. Value: 1. Indicates support for MSRC001_020[A,8,6,4,2,0] and MSRC001_020[B,9,7,5,3,1] .
22	TopologyExtensions: topology extensions support. Value: IF (CPUID Fn8000_0001_EBX[Pkg-Type]==0010b) THEN 1 ELSE 0 ENDIF. Indicates support for CPUID Fn8000_001D_EAX_x0-CPUID Fn8000_001E_EDX .
21	TBM: trailing bit manipulation instruction support. Value: 1.
20	Reserved.
19	NodeId. Value: 1. Asserted by BIOS recommendation to MSRC001_1005 . Indicates support for MSRC001_100C[NodeId, NodesPerProcessor] .
18	Reserved.
17	TCE: translation cache extension. Value: MSRC000_0080[TCE] .
16	FMA4: 4-operand FMA instruction support. Value: 1.
15	LWP: lightweight profiling support. Value: 1.
14	Reserved.
13	WDT: watchdog timer support. Value: 1.
12	SKINIT: SKINIT and STGI support. Value: 1.
11	XOP: extended operation support. Value: 1.
10	IBS: Instruction Based Sampling. Value: 1.
9	OSVW: OS Visible Work-around support. Value: 1.
8	3DNowPrefetch: Prefetch and PrefetchW instructions. Value: 1.
7	MisAlignSse: Misaligned SSE Mode. Value: 1.
6	SSE4A: EXTRQ, INSERTQ, MOVNTSS, and MOVNTSD instruction support. Value: 1.
5	ABM: advanced bit manipulation. Value: 1.
4	AltMovCr8: LOCK MOV CR0 means MOV CR8. Value: 1.

Bits	Description
3	ExtApicSpace: extended APIC register space. Value: 1.
2	SVM: Secure Virtual Mode feature. Value: Product-specific. Indicates support for: VMRUN, VMLOAD, VMSAVE, CLGI, VMMCALL, and INVLPGA.
1	CmpLegacy: core multi-processing legacy mode. Value: Product-specific. 1=Multi core product (CPUID Fn8000_0008_ECX[NC] != 0). 0=Single core product (CPUID Fn8000_0008_ECX[NC] == 0).
0	LahfSahf: LAHF/SAHF instructions. Value: 1.

CPUID Fn8000_0001_EDX Feature Identifiers

These values can be over-written by [MSRC001_1005](#).

Bits	Description
31	3DNow: 3DNow!™ instructions. Value: 0.
30	3DNowExt: AMD extensions to 3DNow!™ instructions. Value: 0.
29	LM: long mode. Value: 1.
28	Reserved.
27	RDTSCP: RDTSCP instruction. Value: 1.
26	Page1GB: one GB large page support. Value: 1.
25	FFXSR: FXSAVE and FXRSTOR instruction optimizations. Value: 1.
24	FXSR: FXSAVE and FXRSTOR instructions. Value: 1.
23	MMX: MMX™ instructions. Value: 1.
22	MmxExt: AMD extensions to MMX™ instructions. Value: 1.
21	Reserved.
20	NX: no-execute page protection. Value: 1.
19:18	Reserved.
17	PSE36: page-size extensions. Value: 1.
16	PAT: page attribute table. Value: 1.
15	CMOV: conditional move instructions, CMOV, FCOMI, FCMOV. Value: 1.
14	MCA: machine check architecture, MCG_CAP. Value: 1.
13	PGE: page global extension, CR4.PGE. Value: 1.
12	MTRR: memory-type range registers. Value: 1.
11	SysCallSysRet: SYSCALL and SYSRET instructions. Value: 1.
10	Reserved.
9	APIC: advanced programmable interrupt controller (APIC) exists and is enabled. Value: MSR0000_001B[ApicEn] .
8	CMPXCHG8B: CMPXCHG8B instruction. Value: 1.
7	MCE: machine check exception, CR4.MCE. Value: 1.
6	PAE: physical-address extensions (PAE). Value: 1.
5	MSR: model-specific registers (MSRs), with RDMSR and WRMSR instructions. Value: 1.

Bits	Description
4	TSC: time stamp counter, RDTSC/RDTSCP instructions, CR4.TSD. Value: 1.
3	PSE: page-size extensions (4 MB pages). Value: 1.
2	DE: debugging extensions, IO breakpoints, CR4.DE. Value: 1.
1	VME: virtual-mode enhancements. Value: 1.
0	FPU: x87 floating point unit on-chip. Value: 1.

CPUID Fn8000_000[4:2]_E[D,C,B,A]X Processor Name String Identifier

Table 205: Value mapping for CPUID Fn8000_000[4:2]_E[D,C,B,A]X Value

Register	Value
CPUID Fn8000_0002_EAX	MSRC001_0030[31:0]
CPUID Fn8000_0002_EBX	MSRC001_0030[63:32]
CPUID Fn8000_0002_ECX	MSRC001_0031[31:0]
CPUID Fn8000_0002_EDX	MSRC001_0031[63:32]
CPUID Fn8000_0003_EAX	MSRC001_0032[31:0]
CPUID Fn8000_0003_EBX	MSRC001_0032[63:32]
CPUID Fn8000_0003_ECX	MSRC001_0033[31:0]
CPUID Fn8000_0003_EDX	MSRC001_0033[63:32]
CPUID Fn8000_0004_EAX	MSRC001_0034[31:0]
CPUID Fn8000_0004_EBX	MSRC001_0034[63:32]
CPUID Fn8000_0004_ECX	MSRC001_0035[31:0]
CPUID Fn8000_0004_EDX	MSRC001_0035[63:32]

Bits	Description
31:0	ProcName: processor name. These return the ASCII string corresponding to the processor name, stored in MSRC001_00[35:30] [Processor Name String].

CPUID Fn8000_0005_EAX L1 TLB 2M/4M Identifiers

This function provides first level TLB characteristics for 2M and 4M pages shared by each core on a compute unit.

Bits	Description
31:24	L1DTlb2and4MAssoc: data TLB associativity for 2 MB and 4 MB pages. Value: FFh. See: CPUID Fn8000_0005_EDX [L1IcAssoc].
23:16	L1DTlb2and4MSize: data TLB number of entries for 2 MB and 4 MB pages. Value: 64. The value returned is for the number of entries available for the 2 MB page size; 4 MB pages require two 2 MB entries, so the number of entries available for the 4 MB page size is one-half the returned value.

15:8	L1ITlb2and4MAssoc: instruction TLB associativity for 2 MB and 4 MB pages. Value: FFh. See: CPUID Fn8000_0005_EDX[L1IcAssoc] .
7:0	L1ITlb2and4MSize: instruction TLB number of entries for 2 MB and 4 MB pages. Value: 24. The value returned is for the number of entries available for the 2 MB page size; 4 MB pages require two 2 MB entries, so the number of entries available for the 4 MB page size is one-half the returned value.

CPUID Fn8000_0005_EBX L1 TLB 4K Identifiers

This function provides first level TLB characteristics for 4K pages shared by each core on a compute unit.

Bits	Description
31:24	L1DTlb4KAssoc: data TLB associativity for 4 KB pages. Value: FFh. See: CPUID Fn8000_0005_EDX[L1IcAssoc] .
23:16	L1DTlb4KSize: data TLB number of entries for 4 KB pages. Value: 64.
15:8	L1ITlb4KAssoc: instruction TLB associativity for 4 KB pages. See: CPUID Fn8000_0005_EDX[L1IcAssoc] . Value: FFh.
7:0	L1ITlb4KSize: instruction TLB number of entries for 4 KB pages. Value: 48.

CPUID Fn8000_0005_ECX L1 Data Cache Identifiers

This function provides first level data cache characteristics for each core.

Bits	Description
31:24	L1DcSize: L1 data cache size in KB. Value: 16.
23:16	L1DcAssoc: L1 data cache associativity. See: CPUID Fn8000_0005_EDX[L1IcAssoc] . Value: 4.
15:8	L1DcLinesPerTag: L1 data cache lines per tag. Value: 1.
7:0	L1DcLineSize: L1 data cache line size in bytes. Value: 64.

CPUID Fn8000_0005_EDX L1 Instruction Cache Identifiers

This function provides first level instruction cache characteristics shared by each core on a compute unit.

Bits	Description										
31:24	L1IcSize: L1 instruction cache size KB. Value: 64.										
23:16	L1IcAssoc: L1 instruction cache associativity. Value: 2. <table border="1" data-bbox="267 1543 1445 1732"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Reserved</td> </tr> <tr> <td>01h</td> <td>1 way (direct mapped)</td> </tr> <tr> <td>FEh-02h</td> <td>Specifies the associativity; e.g., 04h would indicate a 4-way associativity.</td> </tr> <tr> <td>FFh</td> <td>Fully associative</td> </tr> </tbody> </table>	Bits	Description	00h	Reserved	01h	1 way (direct mapped)	FEh-02h	Specifies the associativity; e.g., 04h would indicate a 4-way associativity.	FFh	Fully associative
Bits	Description										
00h	Reserved										
01h	1 way (direct mapped)										
FEh-02h	Specifies the associativity; e.g., 04h would indicate a 4-way associativity.										
FFh	Fully associative										
15:8	L1IcLinesPerTag: L1 instruction cache lines per tag. Value: 1.										
7:0	L1IcLineSize: L1 instruction cache line size in bytes. Value: 64.										

CPUID Fn8000_0006_EAX L2 TLB 2M/4M Identifiers

This function provides second level TLB characteristics for 2M and 4M pages shared by each core on a com-

pute unit.

Bits	Description
31:28	L2DTlb2and4MAssoc: L2 data TLB associativity for 2 MB and 4 MB pages. Value: 6. See: CPUID Fn8000_0006_ECX[L2Assoc] .
27:16	L2DTlb2and4MSize: L2 data TLB number of entries for 2 MB and 4 MB pages. Value: 1024. The value returned is for the number of entries available for the 2 MB page size; 4 MB pages require two 2 MB entries, so the number of entries available for the 4 MB page size is one-half the returned value.
15:12	L2ITlb2and4MAssoc: L2 instruction TLB associativity for 2 MB and 4 MB pages. Value: 6. See: CPUID Fn8000_0006_ECX[L2Assoc] .
11:0	L2ITlb2and4MSize: L2 instruction TLB number of entries for 2 MB and 4 MB pages. Value: 1024. The value returned is for the number of entries available for the 2 MB page size; 4 MB pages require two 2 MB entries, so the number of entries available for the 4 MB page size is one-half the returned value.

CPUID Fn8000_0006_EBX L2 TLB 4K Identifiers

This function provides second level TLB characteristics for 4K pages shared by each core on a compute unit

Bits	Description
31:28	L2DTlb4KAssoc: L2 data TLB associativity for 4 KB pages. Value: 6. See: CPUID Fn8000_0006_ECX[L2Assoc] .
27:16	L2DTlb4KSize: L2 data TLB number of entries for 4 KB pages. Value: 1024.
15:12	L2ITlb4KAssoc: L2 instruction TLB associativity for 4 KB pages. Value: 4. See: CPUID Fn8000_0006_ECX[L2Assoc] .
11:0	L2ITlb4KSize: L2 instruction TLB number of entries for 4 KB pages. Value: 512.

CPUID Fn8000_0006_ECX L2 Cache Identifiers

This function provides second level cache characteristics shared by each core on a compute unit

Bits	Description												
31:16	L2Size: L2 cache size in KB. Value: Product-specific. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>03FFh-0000h</td> <td>Reserved</td> </tr> <tr> <td>0400h</td> <td>1 MB</td> </tr> <tr> <td>07FFh-0401h</td> <td>Reserved</td> </tr> <tr> <td>0800h</td> <td>2 MB</td> </tr> <tr> <td>FFFFh-0801h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	03FFh-0000h	Reserved	0400h	1 MB	07FFh-0401h	Reserved	0800h	2 MB	FFFFh-0801h	Reserved
<u>Bits</u>	<u>Description</u>												
03FFh-0000h	Reserved												
0400h	1 MB												
07FFh-0401h	Reserved												
0800h	2 MB												
FFFFh-0801h	Reserved												

15:12	L2Assoc: L2 cache associativity. Value: 8.			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	0000b	Disabled.	1000b	16 ways
	0001b	1 way (direct mapped)	1001b	Reserved
	0010b	2 ways	1010b	32 ways
	0011b	Reserved	1011b	48 ways
	0100b	4 ways	1100b	64 ways
	0101b	Reserved	1101b	96 ways
	0110b	8 ways	1110b	128 ways
	0111b	Reserved	1111b	Fully associative
11:8	L2LinesPerTag: L2 cache lines per tag. Value: 1.			
7:0	L2LineSize: L2 cache line size in bytes. Value: 64.			

CPUID Fn8000_0006_EDX L3 Cache Identifiers

This provides the processor's third level cache characteristics shared by all cores.

Bits	Description
31:18	L3Size: L3 cache size. Value: 0.
17:16	Reserved.
15:12	L3Assoc: L3 cache associativity. Value: 0.
11:8	L3LinesPerTag: L3 cache lines per tag. Value: 0.
7:0	L3LineSize: L3 cache line size in bytes. Value: 0.

CPUID Fn8000_0007_E[C,B,A]X Advanced Power Management Information

Bits	Description
31:0	Reserved.

CPUID Fn8000_0007_EDX Advanced Power Management Information

This function provides advanced power management feature identifiers.

Bits	Description
31:11	Reserved.
10	EffFreqRO: read-only effective frequency interface. Value: 1. Indicates presence of MSRC000_00E7 [Read-Only Max Performance Frequency Clock Count (MPerfReadOnly)] and MSRC000_00E8 [Read-Only Actual Performance Frequency Clock Count (APerfReadOnly)].
9	CPB: core performance boost. Value: Product-specific. 1= Indicates presence of MSRC001_0015 [CpbDis] and support for core performance boost. See 2.5.3.1.1 [Application Power Management (APM)].
8	TscInvariant: TSC invariant. Value: 1. The TSC rate is invariant.
7	HwpPstate: hardware P-state control. Value: 1. MSRC001_0061 [P-state Current Limit], MSRC001_0062 [P-state Control] and MSRC001_0063 [P-state Status] exist.
6	100MHzSteps: 100 MHz multiplier Control. Value: 1.

5	Reserved.
4	TM: hardware thermal control (HTC). Value: Product-specific.
3	TTP: THERMTRIP. Value: 1.
2	VID: Voltage ID control. Value: 0. Function replaced by HwPstate.
1	FID: Frequency ID control. Value: 0. Function replaced by HwPstate.
0	TS: Temperature sensor. Value: 1.

CPUID Fn8000_0008_EAX Long Mode Address Size Identifiers

This provides information about the maximum physical and linear address width supported by the processor.

Bits	Description
31:24	Reserved.
23:16	GuestPhysAddrSize: maximum guest physical byte address size in bits. Value: 0. 0=The maximum guest physical address size defined by PhysAddrSize.
15:8	LinAddrSize: Maximum linear byte address size in bits. Value: IF (CPUID Fn8000_0001_EDX[LM]) THEN 30h ELSE 20h ENDIF.
7:0	PhysAddrSize: Maximum physical byte address size in bits. Value: 30h.

CPUID Fn8000_0008_EBX Reserved

Bits	Description
31:0	Reserved.

CPUID Fn8000_0008_ECX APIC ID Size and Core Count

This provides information about the number of cores supported by the processor.

Bits	Description
31:16	Reserved.
15:12	ApicIdCoreIdSize: APIC ID size. Value: 4h. The number of bits in the initial APIC20[ApicId] value that indicate core ID within a processor.
11:8	Reserved.
7:0	NC: number of physical cores - 1. The number of cores in the processor is NC+1 (e.g., if NC=0, then there is one core). See 2.4.3 [Processor Cores and Downcoring] and D18F3x190[DisCore]. Value: Product-specific.

CPUID Fn8000_0008_EDX Reserved

Bits	Description
31:0	Reserved.

CPUID Fn8000_0009 Reserved

Bits	Description
31:0	Reserved.

CPUID Fn8000_000A_EAX SVM Revision

This provides SVM revision. If [CPUID Fn8000_0001_ECX\[SVM\]=0](#) then [CPUID Fn8000_000A_EAX](#) is reserved.

Bits	Description
31:8	Reserved.
7:0	SvmRev: SVM revision. Value: 01h.

CPUID Fn8000_000A_EBX SVM Revision and Feature Identification

This provides SVM revision and feature information. If [CPUID Fn8000_0001_ECX\[SVM\]=0](#) then [CPUID Fn8000_000A_EBX](#) is reserved.

Bits	Description
31:0	NASID: number of address space identifiers (ASID). Value: 10000h.

CPUID Fn8000_000A_ECX Reserved

Bits	Description
31:0	Reserved.

CPUID Fn8000_000A_EDX SVM Feature Identification

This provides SVM feature information. If [CPUID Fn8000_0001_ECX\[SVM\]=0](#) then [CPUID Fn8000_000A_EDX](#) is reserved.

Bits	Description
31:13	Reserved.
12	PauseFilterThreshold: PAUSE filter threshold. Value: 1.
11	Reserved.
10	PauseFilter: pause intercept filter. Value: 1.
9:8	Reserved.
7	DecodeAssists: decode assists. Value: 1.
6	FlushByAsid: flush by ASID. Value: 1.
5	VmcbClean: VMCB clean bits. Value: 1.
4	TscRateMsr: MSR based TSC rate control. Value: 1. 1=Indicates support for TSC ratio MSRC000_0104 .
3	NRIPS: NRIP Save. Value: 1.

Bits	Description
2	SVML: SVM lock. Value: 1.
1	LbrVirt: LBR virtualization. Value: 1.
0	NP: nested paging. Value: 1.

CPUID Fn8000_00[18:0B] Reserved

Bits	Description
31:0	Reserved.

CPUID Fn8000_0019_EAX L1 TLB 1G Identifiers

This function provides 1 GB paging information. The associativity fields are defined by [CPUID Fn8000_0006_ECX\[L2Assoc\]](#).

Bits	Description
31:28	L1DTlb1GAssoc: L1 data TLB associativity for 1 GB pages. Value: Fh.
27:16	L1DTlb1GSize: L1 data TLB number of entries for 1 GB pages. Value: 64.
15:12	L1ITlb1GAssoc: L1 instruction TLB associativity for 1 GB pages. Value: Fh.
11:0	L1ITlb1GSize: L1 instruction TLB number of entries for 1 GB pages. Value: 24.

CPUID Fn8000_0019_EBX L2 TLB 1G Identifiers

This function provides 1 GB paging information. The associativity fields are defined by [CPUID Fn8000_0006_ECX\[L2Assoc\]](#).

Bits	Description
31:28	L2DTlb1GAssoc: L2 data TLB associativity for 1 GB pages. Value: 6.
27:16	L2DTlb1GSize: L2 data TLB number of entries for 1 GB pages. Value: 1024.
15:12	L2ITlb1GAssoc: L2 instruction TLB associativity for 1 GB pages. Value: 6.
11:0	L2ITlb1GSize: L2 instruction TLB number of entries for 1 GB pages. Value: 1024.

CPUID Fn8000_0019_E[D,C]X Reserved

Bits	Description
31:0	Reserved.

CPUID Fn8000_001A_EAX Performance Optimization Identifiers

This function returns performance related information. For more details on how to use these bits to optimize software, see the optimization guide.

Bits	Description
31:2	Reserved.
1	MOVU: movu. Value: 1.
0	FP128: fp128. Value: 1.

CPUID Fn8000_001A_E[D,C,B]X Reserved

Bits	Description
31:0	Reserved.

CPUID Fn8000_001B_EAX Instruction Based Sampling Identifiers

This function returns IBS feature information.

Bits	Description
31:8	Reserved.
7	RipInvalidChk: invalid RIP indication supported. Value: 1.
6	OpCntExt: IbsOpCurCnt and IbsOpMaxCnt extend by 7 bits. Value: 1.
5	BrnTrgt: branch target address reporting supported. Value: 1.
4	OpCnt: op counting mode supported. Value: 1.
3	RdWrOpCnt: read write of op counter supported. Value: 1.
2	OpSam: IBS execution sampling supported. Value: 1.
1	FetchSam: IBS fetch sampling supported. Value: 1.
0	IBSFFV: IBS feature flags valid. Value: 1.

CPUID Fn8000_001B_E[D,C,B]X Instruction Based Sampling Identifiers

Bits	Description
31:0	Reserved.

CPUID Fn8000_001C_EAX Lightweight Profiling Capabilities 0

This function returns IBS feature information; see the Lightweight Profiling Specification section titled “Detecting LWP”. If [CPUID Fn8000_0001_ECX\[LWP\]](#) = 0 then CPUID Fn8000_001C_E[D,C,B,A]X is reserved.

Bits	Description
31	LwpInt: interrupt on threshold overflow available. Value: MSRC000_0105[LwpInt] . 1=Interrupt on threshold overflow is available.
30	LwpPTSC: performance time stamp counter in event record. Value: 0. 1=Enable storing performance time stamp in event record.
29	LwpCont: sampling in continuous mode. Value: 0. 1=Enable continuous mode. 0=Enable synchronized mode.

Bits	Description
28:7	Reserved.
6	LwpRNH: core reference clocks not halted event available. Value: MSRC000_0105[LwpRNH] . 1=Core reference clocks not halted event is available.
5	LwpCNH: core clocks not halted event available. Value: MSRC000_0105[LwpCNH] . 1=Core clocks not halted event is available.
4	LwpDME: DC miss event available. Value: MSRC000_0105[LwpDME] . 1=DC miss event is available.
3	LwpBRE: branch retired event available. Value: MSRC000_0105[LwpBRE] . 1=Branch retired event is available.
2	LwpIRE: instructions retired event available. Value: MSRC000_0105[LwpIRE] . 1=Instructions retired event is available.
1	LwpVAL: LWPVAL instruction available. Value: MSRC000_0105[LwpVAL] . 1=LWPVAL instruction is available.
0	LwpAvail: LWP available. Value: XCR0[62]. 1=LWP is available.

CPUID Fn8000_001C_EBX Lightweight Profiling Capabilities 0

See [CPUID Fn8000_001C_EAX](#).

Bits	Description
31:24	LwpEventOffset: offset to the EventInterval1 field. Value: . Offset from the start of the LWPCB to the EventInterval1 field.
23:16	LwpMaxEvents: maximum EventId. Value: . Maximum EventId value that is supported.
15:8	LwpEventSize: event record size. Value: 20h. Size in bytes of an event record in the LWP event ring buffer.
7:0	LwpCbSize: control block size. Value: 13h. Size in quadwords of the LWPCB.

CPUID Fn8000_001C_ECX Lightweight Profiling Capabilities 0

See [CPUID Fn8000_001C_EAX](#).

Bits	Description
31	LwpCacheLatency: cache latency filtering supported. Value: 0. 1=Cache-related events can be filtered by latency.
30	LwpCacheLevels: cache level filtering supported. Value: 0. 1=Cache-related events can be filtered by the cache level that returned the data.
29	LwpIpFiltering: IP filtering supported. Value: 0. 1=IP filtering is supported.
28	LwpBranchPrediction: branch prediction filtering supported. Value: 0. 1=Branches Retired events can be filtered based on whether the branch was predicted properly.
27:24	Reserved.
23:16	LwpMinBufferSize: event ring buffer size. Value: 01h. Minimum size of the LWP event ring buffer, in units of 32 event records.
15:9	LwpVersion: version. Value: 0000001b. Version of LWP implementation.

Bits	Description
8:6	LwpLatencyRnd: amount cache latency is rounded. Value: 0. The amount by which cache latency is rounded.
5	LwpDataAddress: data cache miss address valid. Value: 0. 1=Address is valid for cache miss event records.
4:0	LwpLatencyMax: latency counter bit size. Value: 0. Size in bits of the cache latency counters.

CPUID Fn8000_001C_EDX Lightweight Profiling Capabilities 0

See [CPUID Fn8000_001C_EAX](#).

Bits	Description
31	LwpInt: interrupt on threshold overflow supported. Value: 1. 1=Interrupt on threshold overflow is supported.
30:7	Reserved.
6	LwpRNH: core reference clocks not halted event supported. Value: 0. 1=Core reference clocks not halted event is supported.
5	LwpCNH: core clocks not halted event supported. Value: 0. 1=Core clocks not halted event is supported.
4	LwpDME: DC miss event supported. Value: 0. 1=DC miss event is supported.
3	LwpBRE: branch retired event supported. Value: 1. 1=Branch retired event is supported.
2	LwpIRE: instructions retired event supported. Value: 1. 1=Instructions retired event is supported.
1	LwpVAL: LWPVAL instruction supported. Value: 1. 1=LWPVAL instruction is supported.
0	LwpAvail: lightweight profiling supported. Value: 1. 1=Lightweight profiling is supported.

CPUID Fn8000_001D_EAX_x0 Cache Properties

[CPUID Fn8000_001D_EAX_x0](#) reports topology information for the DC.

If (CPUID Fn8000_0001_ECX[TopologyExtensions]==0) then CPUID Fn8000_001D_E[D,C,B,A]X is reserved.

Table 206: ECX mapping to Cache Type for CPUID Fn8000_001D_E[D,C,B,A]X

ECX	Cache Type
0	DC
1	IC
2	L2
3	Null

Bits	Description
31:26	Reserved.
25:14	NumSharingCache: number of cores sharing cache. Value: 000h. The number of cores sharing this cache is NumSharingCache+1.
13:10	Reserved.

Bits	Description												
9	FullyAssociative: fully associative cache. Value: 0. 1=Cache is fully associative.												
8	SelfInitialization: cache is self-initializing. Value: 1. 1=Cache is self initializing; cache does not need software initialization.												
7:5	CacheLevel: cache level. Identifies the cache level. Value: 001b. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Reserved.</td> </tr> <tr> <td>001b</td> <td>Level 1</td> </tr> <tr> <td>010b</td> <td>Level 2</td> </tr> <tr> <td>011b</td> <td>Level 3</td> </tr> <tr> <td>111b-100b</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	000b	Reserved.	001b	Level 1	010b	Level 2	011b	Level 3	111b-100b	Reserved.
Bits	Description												
000b	Reserved.												
001b	Level 1												
010b	Level 2												
011b	Level 3												
111b-100b	Reserved.												
4:0	CacheType: cache type. Identifies the type of cache. Value: 01h. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Null; no more caches.</td> </tr> <tr> <td>01h</td> <td>Data cache</td> </tr> <tr> <td>02h</td> <td>Instruction cache</td> </tr> <tr> <td>03h</td> <td>Unified cache</td> </tr> <tr> <td>1Fh-04h</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	00h	Null; no more caches.	01h	Data cache	02h	Instruction cache	03h	Unified cache	1Fh-04h	Reserved.
Bits	Description												
00h	Null; no more caches.												
01h	Data cache												
02h	Instruction cache												
03h	Unified cache												
1Fh-04h	Reserved.												

CPUID Fn8000_001D_EAX_x1 Cache Properties

CPUID Fn8000_001D_EAX_x1 reports topology information for the IC. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:26	Reserved.
25:14	NumSharingCache: number of cores sharing cache. Value: IF (D18F5x80[<i>DualCore</i> [0]]==1) THEN 001h ELSE 000h ENDIF. See: CPUID Fn8000_001D_EAX_x0[NumSharingCache].
13:10	Reserved.
9	FullyAssociative: fully associative cache. Value: 0. See: CPUID Fn8000_001D_EAX_x0[FullyAssociative].
8	SelfInitialization: cache is self-initializing. Value: 1. See: CPUID Fn8000_001D_EAX_x0[SelfInitialization].
7:5	CacheLevel: cache level. Identifies the cache level. Value: 001b. See: CPUID Fn8000_001D_EAX_x0[CacheLevel].
4:0	CacheType: cache type. Value: 02h. See: CPUID Fn8000_001D_EAX_x0[CacheType].

CPUID Fn8000_001D_EAX_x2 Cache Properties

CPUID Fn8000_001D_EAX_x2 reports topology information for the L2. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:26	Reserved.
25:14	NumSharingCache: number of cores sharing cache. Value:IF (D18F5x80[<i>DualCore</i> [0]]==1) THEN 001h ELSE 000h ENDIF. See: CPUID Fn8000_001D_EAX_x0[NumSharingCache].

Bits	Description
13:10	Reserved.
9	FullyAssociative: fully associative cache. Value: 0. See: CPUID Fn8000_001D_EAX_x0[FullyAssociative] .
8	SelfInitialization: cache is self-initializing. Value: 1. See: CPUID Fn8000_001D_EAX_x0[SelfInitialization] .
7:5	CacheLevel: cache level. Identifies the cache level. Value: 010b. See: CPUID Fn8000_001D_EAX_x0[CacheLevel] .
4:0	CacheType: cache type. Value: 03h. See: CPUID Fn8000_001D_EAX_x0[CacheType] .

CPUID Fn8000_001D_EAX_x3 Cache Properties

See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:26	Reserved.
25:14	NumSharingCache: number of cores sharing cache. Value: 0.
13:10	Reserved.
9	FullyAssociative: fully associative cache. Value: 0. See: CPUID Fn8000_001D_EAX_x0[FullyAssociative] .
8	SelfInitialization: cache is self-initializing. Value: 0. See: CPUID Fn8000_001D_EAX_x0[SelfInitialization] .
7:5	CacheLevel: cache level. Identifies the cache level. Value: 000b. See: CPUID Fn8000_001D_EAX_x0[CacheLevel] .
4:0	CacheType: cache type. Value: 00h. See: CPUID Fn8000_001D_EAX_x0[CacheType] .

CPUID Fn8000_001D_EBX_x0 Cache Properties

[CPUID Fn8000_001D_EBX_x0](#) reports topology information for the DC. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:22	CacheNumWays: cache number of ways. Value: 003h. Cache number of ways is CacheNumWays+1.
21:12	CachePhysPartitions: cache physical line partitions. Value: 000h. Cache partitions is CachePhysPartitions+1.
11:0	CacheLineSize: cache line size in bytes. Value: 03Fh. Cache line size in bytes is CacheLineSize+1.

CPUID Fn8000_001D_EBX_x1 Cache Properties

[CPUID Fn8000_001D_EBX_x1](#) reports topology information for the IC. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:22	CacheNumWays: cache number of ways. Value: 001h. See: CPUID Fn8000_001D_EBX_x0 [CacheNumWays].
21:12	CachePhysPartitions: cache physical line partitions. Value: 000h. See: CPUID Fn8000_001D_EBX_x0 [CachePhysPartitions].
11:0	CacheLineSize: cache line size in bytes. Value: 03Fh. See: CPUID Fn8000_001D_EBX_x0 [CacheLineSize].

CPUID Fn8000_001D_EBX_x2 Cache Properties

[CPUID Fn8000_001D_EBX_x2](#) reports topology information for the L2. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:22	CacheNumWays: cache number of ways. Value: 00Fh. See: CPUID Fn8000_001D_EBX_x0 [CacheNumWays].
21:12	CachePhysPartitions: cache physical line partitions. Value: 000h. See: CPUID Fn8000_001D_EBX_x0 [CachePhysPartitions].
11:0	CacheLineSize: cache line size in bytes. Value: 03Fh. See: CPUID Fn8000_001D_EBX_x0 [CacheLineSize].

CPUID Fn8000_001D_EBX_x3 Cache Properties

See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:0	Reserved.

CPUID Fn8000_001D_ECX_x0 Cache Properties

[CPUID Fn8000_001D_ECX_x0](#) reports topology information for the DC. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:0	CacheNumSets: cache number of sets. Value: 0000_003Fh. Cache number of sets is CacheNumSets+1.

CPUID Fn8000_001D_ECX_x1 Cache Properties

[CPUID Fn8000_001D_ECX_x1](#) reports topology information for the IC. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:0	CacheNumSets: cache number of sets. Value: 0000_01FFh. See: CPUID Fn8000_001D_ECX_x0 [CacheNumSets].

CPUID Fn8000_001D_ECX_x2 Cache Properties

CPUID Fn8000_001D_ECX_x2 reports topology information for the L2. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:0	CacheNumSets: cache number of sets. Value: Product-specific. See: CPUID Fn8000_001D_ECX_x0[CacheNumSets].

CPUID Fn8000_001D_ECX_x3 Cache Properties

See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:0	Reserved.

CPUID Fn8000_001D_EDX_x0 Cache Properties

CPUID Fn8000_001D_EDX_x0 reports topology information for the DC. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:2	Reserved.
1	CacheInclusive: cache inclusive. Value: 0. 0=Cache is not inclusive of lower cache levels. 1=Cache is inclusive of lower cache levels.
0	WBINVD: Write-Back Invalidate/Invalidate. Value: 0. 0=WBINVD/INVD invalidates all lower level caches of non-originating cores sharing this cache. 1=WBINVD/INVD not guaranteed to invalidate all lower level caches of non-originating cores sharing this cache.

CPUID Fn8000_001D_EDX_x1 Cache Properties

CPUID Fn8000_001D_EDX_x1 reports topology information for the IC. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:2	Reserved.
1	CacheInclusive: cache inclusive. Value: 0. See: CPUID Fn8000_001D_EDX_x0[CacheInclusive].
0	WBINVD: Write-Back Invalidate/Invalidate. Value: 0. See: CPUID Fn8000_001D_EDX_x0[WBINVD].

CPUID Fn8000_001D_EDX_x2 Cache Properties

CPUID Fn8000_001D_EDX_x2 reports topology information for the L2. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:2	Reserved.
1	CacheInclusive: cache inclusive. Value: 0. See: CPUID Fn8000_001D_EDX_x0[CacheInclusive] .
0	WBINVD: Write-Back Invalidate/Invalidate. Value: 1. See: CPUID Fn8000_001D_EDX_x0[WBINVD] .

CPUID Fn8000_001D_EDX_x3 Cache Properties

See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:0	Reserved.

CPUID Fn8000_001E_EAX Extended APIC ID

If [CPUID Fn8000_0001_ECX\[TopologyExtensions\]](#)==0 then [CPUID Fn8000_001E_E\[D,C,B,A\]X](#) is reserved. If ([MSR0000_001B\[ApicEn\]](#)==0) then [CPUID Fn8000_001E_EAX\[ExtendedApicId\]](#) is reserved.

Bits	Description
31:0	ExtendedApicId: extended APIC ID. Value: {000000h, APIC20[31:24] }.

CPUID Fn8000_001E_EBX Compute Unit Identifiers

See [CPUID Fn8000_001E_EAX](#).

Bits	Description
31:10	Reserved.
9:8	CoresPerComputeUnit: cores per compute unit. Value: Product-specific. The number of cores per compute unit is CoresPerComputeUnit +1.
7:0	ComputeUnitId: compute unit ID. Value: Product-specific. Identifies the processor compute unit ID.

CPUID Fn8000_001E_ECX Node Identifiers

See [CPUID Fn8000_001E_EAX](#).

Bits	Description
31:11	Reserved.
10:8	NodesPerProcessor. Value: 0.
7:0	NodeId. Value: 0.

CPUID Fn8000_001E_EDX Reserved

See [CPUID Fn8000_001E_EAX](#).

Bits	Description
31:0	Reserved.

3.18 MSRs - MSR0000_XXXX

See 3.1 [Register Descriptions and Mnemonics] for a description of the register naming convention. MSRs are accessed through x86 WRMSR and RDMSR instructions.

MSR0000_0000 Load-Store MCA Address

Bits	Description
63:0	Alias of MSR0000_0402 .

MSR0000_0001 Load-Store MCA Status

Bits	Description
63:0	Alias of MSR0000_0401 .

MSR0000_0010 Time Stamp Counter (TSC)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:32	TSC[63:32]: time stamp counter high. See: TSC[31:0].
31:0	TSC[31:0]: time stamp counter low. Read-write; updated-by-hardware. TSC[63:0] = {TSC[63:32], TSC[31:0]}. The TSC increments at the P0 frequency. This field uses software P-state numbering. See 2.5.3.1.2.1 [Software P-state Numbering] . The TSC counts at the same rate in all P-states, all C states, S0, or S1. The TSC rate is affected by MSRC000_0104 [Time Stamp Counter Ratio (TscRateMsr)] . The value (TSC/TSCRatio) is the TSC P0 frequency based value (as if TSCRatio==1.0) when (TSCRatio!=1.0).

MSR0000_001B APIC Base Address (APIC_BAR)

Bits	Description
63:48	MBZ.
47:12	ApicBar[47:12]: APIC base address register. Read-write. Reset: 0000FEE00h. Specifies the base address, physical address [47:12], for the APICXX register set in xAPIC mode. See 2.4.8.1.2 [APIC Register Space] .
11	ApicEn: APIC enable. Read-write. Reset: 0. See 2.4.8.1.2 [APIC Register Space] . 1=Local APIC is enabled in xAPIC mode.
10:9	MBZ.
8	BSC: boot strap core. Read-write; updated-by-hardware. Reset: x. 1=The core is the boot core of the BSP. 0=The core is not the boot core of the BSP.
7:0	MBZ.

MSR0000_002A Cluster ID (EBL_CR_POWERON)

Reset: 0. Read; GP-write. Writes to this register result in a GP faults with error code 0.

Bits	Description
63:18	MBZ.
17:16	ClusterID . This field does not affect hardware.
15:0	MBZ.

MSR0000_00E7 Max Performance Frequency Clock Count (MPERF)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	MPERF: maximum core clocks counter . Read-write; Updated-by-hardware . Incremented by hardware at the P0 frequency while the core is in C0. This register does not increment when the core is in the stop-grant state. In combination with MSR0000_00E8 , this is used to determine the effective frequency of the core. This field uses software P-state numbering. See MSRC001_0015 [EffFreqCntMwait], 2.5.3.3 [Effective Frequency], and 2.5.3.1.2.1 [Software P-state Numbering]. This register is not affected by writes to MSR0000_00E7 .

MSR0000_00E8 Actual Performance Frequency Clock Count (APERF)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	APERF: actual core clocks counter . Read-write; Updated-by-hardware . This register increments in proportion to the actual number of core clocks cycles while the core is in C0. The register does not increment when the core is in the stop-grant state. See MSR0000_00E7 . This register is not affected by writes to MSR0000_00E8 .

MSR0000_00FE MTRR Capabilities (MTRRcap)

Read; GP-write. Reset: 0000_0000_0000_0508h.

Bits	Description
63:11	Reserved.
10	MtrrCapWc: write-combining memory type . 1=The write combining memory type is supported.
9	Reserved.
8	MtrrCapFix: fixed range register . 1=Fixed MTRRs are supported.
7:0	MtrrCapVCnt: variable range registers count . Specifies the number of variable MTRRs supported.

MSR0000_0174 SYSENTER CS (SYSENTER_CS)

Bits	Description
63:32	RAZ.
31:16	Reserved.
15:0	SysEnterCS: SYSENTER target CS . Read-write. Reset: 0. Holds the called procedure code segment.

MSR0000_0175 SYSENTER ESP (SYSENTER_ESP)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:32	Reserved.
31:0	SysEnterESP: SYSENTER target SP. Read-write. Holds the called procedure stack pointer.

MSR0000_0176 SYSENTER EIP (SYSENTER_EIP)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:32	Reserved.
31:0	SysEnterEIP: SYSENTER target IP. Read-write. Holds the called procedure instruction pointer.

MSR0000_0179 Global Machine Check Capabilities (MCG_CAP)

Read; GP-write.

Bits	Description
63:9	Reserved.
8	McgCtlP: MCG_CTL register present. Value: 1. 1=The machine check control registers (MCI_CTL) are present. See 2.15.1 [Machine Check Architecture]
7:0	Count. Value: 07h. Indicates the number of error reporting banks visible to each core.

MSR0000_017A Global Machine Check Status (MCG_STAT)

Reset: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture].

Bits	Description
63:3	Reserved.
2	MCIP: machine check in progress. Read-write; set-by-hardware. 1=A machine check is in progress.
1	EIPV: error instruction pointer valid. Read-write; Updated-by-hardware. 1=The instruction pointer that was pushed onto the stack by the machine check mechanism references the instruction that caused the machine check error.
0	RIPV: restart instruction pointer valid. Read-write; Updated-by-hardware. 1=Program execution can be reliably restarted at the EIP address on the stack. 0=The interrupt was not precise and/or the process (task) context may be corrupt; continued operation of this process may not be possible without intervention, however system processing or other processes may be able to continue with appropriate software clean up.

MSR0000_017B Global Machine Check Exception Reporting Control (MCG_CTL)

Reset: 0000_0000_0000_0000h. This registers controls enablement of the individual error reporting banks; see 2.15.1 [Machine Check Architecture]. When a machine check register bank is not enabled in MCG_CTL, errors for that bank are not logged or reported, and actions enabled through the MCA are not taken; each

MCi_CTL register identifies which errors are still corrected when MCG_CTL[i] is disabled.

Bits	Description
63:7	Unused.
6	MC6En: MC6 register bank enable. Read-write. 1=The FP machine check register bank is enabled.
5	MC5En: MC5 register bank enable. Read-write. 1=The EX machine check register bank is enabled.
4	MC4En: MC4 register bank enable. Read-write. 1=The NB machine check register bank is enabled for all cores of the node.
3	Unused.
2	MC2En: MC2 register bank enable. Read-write. 1=The CU machine check register bank is enabled.
1	MC1En: MC1 register bank enable. Read-write. 1=The IF machine check register bank is enabled.
0	MC0En: MC0 register bank enable. Read-write. 1=The LS machine check register bank is enabled.

MSR0000_01D9 Debug Control (DBG_CTL_MSR)

Bits	Description
63:7	Reserved.
6	MBZ.
5:2	PB: performance monitor pin control. Read-write. Reset: 0. This field does not control any hardware.
1	BTF. Read-write. Reset: 0. 1=Enable branch single step.
0	LBR. Read-write. Reset: 0. 1=Enable last branch record.

MSR0000_01DB Last Branch From IP (BR_FROM)

Reset: 0000_0000_0000_0000h. Read; GP-write; [Not-same-for-all](#).

Bits	Description
63:0	LastBranchFromIP. Loaded with the segment offset of the branch instruction.

MSR0000_01DC Last Branch To IP (BR_TO)

Reset: 0000_0000_0000_0000h. Read; GP-write; [Not-same-for-all](#).

Bits	Description
63:0	LastBranchToIP. Holds the target RIP of the last branch that occurred before an exception or interrupt.

MSR0000_01DD Last Exception From IP

Reset: 0000_0000_0000_0000h. Read; GP-write; [Not-same-for-all](#).

Bits	Description
63:0	LastIntFromIP. Holds the source RIP of the last branch that occurred before the exception or interrupt.

MSR0000_01DE Last Exception To IP

Reset: 0000_0000_0000_0000h. Read; GP-write; **Not-same-for-all**.

Bits	Description
63:0	LastIntToIP . Holds the target RIP of the last branch that occurred before the exception or interrupt.

MSR0000_020[F:0] Variable-Size MTRRs Base/Mask

Each MTRR ([MSR0000_020\[F:0\] \[Variable-Size MTRRs Base/Mask\]](#), [MSR0000_02\[6F:68,59:58,50\]](#), or [MSR0000_02FF \[MTRR Default Memory Type \(MTRRdefType\)\]](#)) specifies a physical address range and a corresponding memory type (MemType) associated with that range. Setting the memory type to an unsupported value results in a #GP.

The variable-size MTRRs come in pairs of base and mask registers (MSR0000_0200 and MSR0000_0201 are the first pair, etc.). Variables MTRRs are enabled through [MSR0000_02FF \[MTRR Default Memory Type \(MTRRdefType\)\] \[MtrrDefTypeEn\]](#). A core access--with address CPUAddr--is determined to be within the address range of a variable-size MTRR if the following equation is true:

$\text{CPUAddr}[47:12] \& \text{PhyMask}[47:12] == \text{PhyBase}[47:12] \& \text{PhyMask}[47:12]$.

For example, if the variable MTRR spans 256 KB and starts at the 1 MB address. The PhyBase would be set to 0_0010_0000h and the PhyMask to F_FFFC_0000h (with zeros filling in for bits[11:0]). This results in a range from 0_0010_0000h to 0_0013_FFFF.

MSR0000_020[E,C,A,8,6,4,2,0] Variable-Size MTRRs Base

SharedC.

Table 207: Register Mapping for MSR0000_020[E,C,A,8,6,4,2,0]

Register	Function
MSR0000_0200	Range 0
MSR0000_0202	Range 1
MSR0000_0204	Range 2
MSR0000_0206	Range 3
MSR0000_0208	Range 4
MSR0000_020A	Range 5
MSR0000_020C	Range 6
MSR0000_020E	Range 7

Table 208: Memory Type Definition

Bits	Description
000b	UC or uncacheable.
001b	WC or write combining.
011b-010b	Reserved
100b	WT or write through.

Table 208: Memory Type Definition

Bits	Description
101b	WP or write protect.
110b	WB or write back.
111b	Reserved

Bits	Description
63:48	MBZ.
47:12	PhyBase: base address. Read-write. Reset: 0.
11:3	MBZ.
2:0	MemType: memory type. Read-write. Reset: 0. Address range from 00000h to 0FFFFh. See: Table 208 [Memory Type Definition] .

MSR0000_020[F,D,B,9,7,5,3,1] Variable-Size MTRRs Mask

SharedC.

Table 209: Register Mapping for MSR0000_020[F,D,B,9,7,5,3,1]

Register	Function
MSR0000_0201	Range 0
MSR0000_0203	Range 1
MSR0000_0205	Range 2
MSR0000_0207	Range 3
MSR0000_0209	Range 4
MSR0000_020B	Range 5
MSR0000_020D	Range 6
MSR0000_020F	Range 7

Bits	Description
63:48	MBZ.
47:12	PhyMask: address mask. Read-write. Reset: 0.
11	Valid: valid. Read-write. Reset: 0. 1=The variable-size MTRR pair is enabled.
10:0	MBZ.

MSR0000_02[6F:68,59:58,50] Fixed-Size MTRRs

SharedC. Reset: 0000_0000_0000_0000h. See [MSR0000_020\[F:0\]](#) for general MTRR information. Fixed MTRRs are enabled through [MSR0000_02FF\[MtrrDefTypeFixEn and MtrrDefTypeEn\]](#). For addresses below 1MB, the appropriate Fixed MTRRs override the default access destination. Each fixed MTRR includes two bits, RdDram and WrDram, that determine the destination based on the access type. See [2.4.5.1.2 \[Determining The Access Destination for Core Accesses\]](#).

Table 210: Register Mapping for MSR0000_02[6F:68,59:58,50]

Register	Function
MSR0000_0250	64K Range
MSR0000_0258	16K_0 Range
MSR0000_0259	16K_1 Range
MSR0000_0268	4K_0 Range
MSR0000_0269	4K_1 Range
MSR0000_026A	4K_2 Range
MSR0000_026B	4K_3 Range
MSR0000_026C	4K_4 Range
MSR0000_026D	4K_5 Range
MSR0000_026E	4K_6 Range
MSR0000_026F	4K_7 Range

Table 211: Fixed-size MTRR size and Range Mapping

Register	Bits							
	63:56	55:48	47:40	39:32	31:24	23:16	15:8	7:0
MSR0000_0250	64K_70000	64K_60000	64K_50000	64K_40000	64K_30000	64K_20000	64K_10000	64K_00000
MSR0000_0258	16K_9C000	16K_98000	16K_94000	16K_90000	16K_8C000	16K_88000	16K_84000	16K_80000
MSR0000_0259	16K_BC000	16K_B8000	16K_B4000	16K_B0000	16K_AC000	16K_A8000	16K_A4000	16K_A0000
MSR0000_0268	4K_C7000	4K_C6000	4K_C5000	4K_C4000	4K_C3000	4K_C2000	4K_C1000	4K_C0000
MSR0000_0269	4K_CF000	4K_CE000	4K_CD000	4K_CC000	4K_CB000	4K_CA000	4K_C9000	4K_C8000
MSR0000_026A	4K_D7000	4K_D6000	4K_D5000	4K_D4000	4K_D3000	4K_D2000	4K_D1000	4K_D0000
MSR0000_026B	4K_DF000	4K_DE000	4K_DD000	4K_DC000	4K_DB000	4K_DA000	4K_D9000	4K_D8000
MSR0000_026C	4K_E7000	4K_E6000	4K_E5000	4K_E4000	4K_E3000	4K_E2000	4K_E1000	4K_E0000
MSR0000_026D	4K_EF000	4K_EE000	4K_ED000	4K_EC000	4K_EB000	4K_EA000	4K_E9000	4K_E8000
MSR0000_026E	4K_F7000	4K_F6000	4K_F5000	4K_F4000	4K_F3000	4K_F2000	4K_F1000	4K_F0000
MSR0000_026F	4K_FF000	4K_FE000	4K_FD000	4K_FC000	4K_FB000	4K_FA000	4K_F9000	4K_F8000

Bits	Description
63:61	MBZ.
60	RdDram: read DRAM. See: MSR0000_02[6F:68,59:58,50][4] .
59	WrDram: write DRAM. See: MSR0000_02[6F:68,59:58,50][3] .
58:56	MemType: memory type. See: MSR0000_02[6F:68,59:58,50][2:0] .
55:53	MBZ.
52	RdDram: read DRAM. See: MSR0000_02[6F:68,59:58,50][4] .
51	WrDram: write DRAM. See: MSR0000_02[6F:68,59:58,50][3] .
50:48	MemType: memory type. See: MSR0000_02[6F:68,59:58,50][2:0] .
47:45	MBZ.
44	RdDram: read DRAM. See: MSR0000_02[6F:68,59:58,50][4] .

43	WrDram: write DRAM. See: MSR0000_02[6F:68,59:58,50][3] .
42:40	MemType: memory type. See: MSR0000_02[6F:68,59:58,50][2:0] .
39:37	MBZ.
36	RdDram: read DRAM. See: MSR0000_02[6F:68,59:58,50][4] .
35	WrDram: write DRAM. See: MSR0000_02[6F:68,59:58,50][3] .
34:32	MemType: memory type. See: MSR0000_02[6F:68,59:58,50][2:0] .
31:29	MBZ.
28	RdDram: read DRAM. See: MSR0000_02[6F:68,59:58,50][4] .
27	WrDram: write DRAM. See: MSR0000_02[6F:68,59:58,50][3] .
26:24	MemType: memory type. See: MSR0000_02[6F:68,59:58,50][2:0] .
23:21	MBZ.
20	RdDram: read DRAM. See: MSR0000_02[6F:68,59:58,50][4] .
19	WrDram: write DRAM. See: MSR0000_02[6F:68,59:58,50][3] .
18:16	MemType: memory type. See: MSR0000_02[6F:68,59:58,50][2:0] .
15:13	MBZ.
12	RdDram: read DRAM. See: MSR0000_02[6F:68,59:58,50][4] .
11	WrDram: write DRAM. See: MSR0000_02[6F:68,59:58,50][3] .
10:8	MemType: memory type. See: MSR0000_02[6F:68,59:58,50][2:0] .
7:5	MBZ.
4	RdDram: read DRAM. Read-write. IF (MSRC001_0010[MtrrFixDramModEn]) THEN Read-write. ELSE MBZ. ENDIF. 0=Read accesses to the range are marked as MMIO. 1=Read accesses to the range are marked as destined for DRAM. Address range from 00000h to 0FFFFh. See: MSRC001_0010[MtrrFixDramEn, MtrrFixDramModEn] .
3	WrDram: write DRAM. Read-write. IF (MSRC001_0010[MtrrFixDramModEn]) THEN Read-write. ELSE MBZ. ENDIF. 0=Write accesses to the range are marked as MMIO. 1=Write accesses to the range are marked as destined for DRAM. Address range from 00000h to 0FFFFh. See: MSRC001_0010[MtrrFixDramEn, MtrrFixDramModEn] .
2:0	MemType: memory type. Read-write. Address range from 00000h to 0FFFFh. See: Table 208 [Memory Type Definition] .

MSR0000_0277 Page Attribute Table (PAT)

This register specifies the memory type based on the PAT, PCD, and PWT bits in the virtual address page tables.

Bits	Description
63:59	MBZ.
58:56	PA7MemType. See: PA0MemType. Reset: 0h. Default UC. MemType for {PAT, PCD, PWT} = 7h.
55:51	MBZ.
50:48	PA6MemType. See: PA0MemType. Reset: 7h. Default UC-. MemType for {PAT, PCD, PWT} = 6h.
47:43	MBZ.
42:40	PA5MemType. See: PA0MemType. Reset: 4h. Default WT. MemType for {PAT, PCD, PWT} = 5h.

39:35	MBZ.																				
34:32	PA4MemType . See: PA0MemType. Reset: 6h. Default WB. MemType for {PAT, PCD, PWT} = 4h.																				
31:27	MBZ.																				
26:24	PA3MemType . See: PA0MemType. Reset: 0h. Default UC. MemType for {PAT, PCD, PWT} = 3h.																				
23:19	MBZ.																				
18:16	PA2MemType . See: PA0MemType. Reset: 7h. Default UC-. MemType for {PAT, PCD, PWT} = 2h.																				
15:11	MBZ.																				
10:8	PA1MemType . See: PA0MemType. Reset: 4h. Default WT. MemType for {PAT, PCD, PWT} = 1h.																				
7:3	MBZ.																				
2:0	PA0MemType . Read-write. Reset: 6h. MemType for {PAT, PCD, PWT} = 0h. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>UC or uncacheable.</td> <td>4h</td> <td>WT or write through.</td> </tr> <tr> <td>1h</td> <td>WC or write combining.</td> <td>5h</td> <td>WP or write protect.</td> </tr> <tr> <td>2h</td> <td>MBZ.</td> <td>6h</td> <td>WB or write back.</td> </tr> <tr> <td>3h</td> <td>MBZ.</td> <td>7h</td> <td>UC- or uncacheable (overridden by WC state).</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	0h	UC or uncacheable.	4h	WT or write through.	1h	WC or write combining.	5h	WP or write protect.	2h	MBZ.	6h	WB or write back.	3h	MBZ.	7h	UC- or uncacheable (overridden by WC state).
Bits	Description	Bits	Description																		
0h	UC or uncacheable.	4h	WT or write through.																		
1h	WC or write combining.	5h	WP or write protect.																		
2h	MBZ.	6h	WB or write back.																		
3h	MBZ.	7h	UC- or uncacheable (overridden by WC state).																		

MSR0000_02FF MTRR Default Memory Type (MTRRdefType)

SharedC. Reset: 0000_0000_0000_0000h. See [MSR0000_020\[F:0\]](#) for general MTRR information.

Bits	Description
63:12	MBZ.
11	MtrrDefTypeEn: variable and fixed MTRR enable . Read-write. 1= MSR0000_020[F:0] [Variable-Size MTRRs Base/Mask], and MSR0000_02[6F:68,59:58,50] [Fixed-Size MTRRs] are enabled. 0=Fixed and variable MTRRs are not enabled.
10	MtrrDefTypeFixEn: fixed MTRR enable . Read-write. 1= MSR0000_02[6F:68,59:58,50] [Fixed-Size MTRRs] are enabled. This field is ignored (and the fixed MTRRs are not enabled) if MSR0000_02FF[MtrrDefTypeEn]=0 .
9:8	MBZ.
7:0	MemType: memory type . Read-write. If MtrrDefTypeEn==1 then MemType specifies the memory type for memory space that is not specified by either the fixed or variable range MTRRs. If MtrrDefTypeEn==0 then the default memory type for all of memory is UC. Valid encodings are {00000b, MSR0000_02[6F:68,59:58,50][2:0] }.

MSR0000_0400 LS Machine Check Control (MC0_CTL)

Read-write. Reset: 0000_0000_0000_0000h. BIOS: FFFF_FFFF_FFFF_FFFFh. BIOS writes MC0_CTL to work around some operating systems that skip MC0_CTL during MCA initialization. See [2.15.1 \[Machine Check Architecture\]](#). See [MSRC001_0044 \[LS Machine Check Control Mask \(MC0_CTL_MASK\)\]](#).

Bits	Description
63:10	Unused.
9	IntErrTyp1: internal error type 1 .
8	IntErrTyp2: internal error type 2 .
7	SRDE: read data errors . System read data errors on cache fill.

6	LFE: line fill error. Uncorrectable error on cache fill.
5	SCBP: SCB parity.
4	SQP: store queue parity.
3	LQP: load queue parity.
2	DatP: data parity.
1	TLBP: TLB parity.
0	TagP: tag parity.

MSR0000_0401 LS Machine Check Status (MC0_STATUS)

Cold reset: 0. See 2.15.1 [Machine Check Architecture]. See MSRC001_0015[McStatusWrEn]. Table 212 describes each error type. Table 213 describes the error codes and status register settings for each error type. MSR0000_0001 is an alias of MSR0000_0401.

Bits	Description
63	Val: valid. Read-write; set-by-hardware. 1=A valid error has been detected (whether it is enabled or not). This bit should be cleared to 0 by software after the register has been read.
62	<p>Overflow: error overflow. Read-write; set-by-hardware. 1=An error was detected while the valid bit (Val) was set; at least one error was not logged. Overflow is set independently of whether the existing error is overwritten.</p> <p>The following hierarchy identifies the error logging priorities.</p> <ol style="list-style-type: none"> 1. Uncorrectable errors 2. Correctable errors <p>The machine check mechanism handles the contents of MCI_STATUS during overflow as follows:</p> <ul style="list-style-type: none"> • Higher priority errors overwrite lower priority errors. • New errors of equal or lower priority do not overwrite existing errors. • Uncorrectable errors which are not logged due to overflow result in setting PCC, unless the new uncorrectable error is of the same type and in the same reportable address range as the existing error.
61	UC: error uncorrected. Read-write; Updated-by-hardware. 1=The error was not corrected by hardware.
60	En: error enable. Read-write; Updated-by-hardware. 1=MCA error reporting is enabled for this error in MCI_CTL.
59	MiscV: miscellaneous error register valid. Read-write. 1=Valid thresholding in MSR0000_0403.
58	AddrV: error address valid. Read-write; Updated-by-hardware. 1=MCI_ADDR contains address information associated with the error.
57	PCC: processor context corrupt. Read-write; Set-by-hardware. 1=Hardware context held by the processor may have been corrupted. Continued operation of the system may have unpredictable results. The error is not recoverable or survivable, and the system should be reinitialized. See 2.15.1.6.1 [MCA Differentiation Between System-Fatal and Process-Fatal Errors].
56:40	Reserved.

39:36	Way: cache way in error. Read-write; Updated-by-hardware. Indicates the cache way in error.												
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Way 0</td> </tr> <tr> <td>1h</td> <td>Way 1</td> </tr> <tr> <td>2h</td> <td>Way 2</td> </tr> <tr> <td>3h</td> <td>Way 3</td> </tr> <tr> <td>Fh-4h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	Way 0	1h	Way 1	2h	Way 2	3h	Way 3	Fh-4h	Reserved
Bits	Description												
0h	Way 0												
1h	Way 1												
2h	Way 2												
3h	Way 3												
Fh-4h	Reserved												
35:21	Reserved.												
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. Logs an extended error code when an error is detected. This model-specific field is used in conjunction with ErrorCode to identify the error sub-type for root cause analysis (see 2.15.1.5 [Error Code]). See Table 213 for values.												
15:0	ErrorCode: error code. Read-write; Updated-by-hardware. See 2.15.1.5 [Error Code].												

Table 212: LS Error Descriptions

Error Type	Error Sub-type	Description ¹	CTL ²	EAC ³
Line Fill Error	-	An uncorrectable error occurred during a line fill from the L2 cache or the NB. (Note: For IO read, may not actually install to L1 cache.)	LineFill-Poison	E
Data Cache Error	Data array	Error occurred in cache data array access.	DatP	D
	SCB	Error occurred in SCB access.	SCBP	D
	STQ	Error occurred in STQ access.	SQP	D
Tag Error	Tag array	A tag error was encountered. If uncorrectable, this errors is system fatal and results in a sync flood.	TagP	D
	STQ	Error occurred in STQ access.	SQP	D
	LDQ	Error occurred in LDQ access.	LQP	D
L1 TLB Error	TLB parity	Parity error in L1 TLB access.	TLBP	D
	TLB multi-match	Lookup hit on multiple entries.		D
	Locked TLB miss	TLB miss occurred after lock granted.		E
System Read Data Error	-	An error occurred during an attempted read of data from the NB. Possible reasons include master abort, target abort.	SRDE	E
Internal Error	Type 1	A condition was detected which prohibits the core from continuing execution.	Internal Error: Type 1	E
	Type 2		Internal Error: Type 2	E

1. CID: core ID. All LS errors are reported to the affected core; see 2.15.1.3.
2. See MSR0000_0400.
3. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See 2.15.1.3.

Table 213: LS Error Signatures

Error Type	Error SubType	Error-CodeExt	Error Code						UC	ADDRV	PCC	Reserved	Reserved			
			Type	PP	T	RRRR	II/TT	LL								
Line Fill Error	-	00001b	MEM	-	-	DRD	D	LG	1	1	0	-	-			
Data Cache Error	Data array	00000b				DRD	D	L1	0/1	1	0	-	-			
	SCB	00011b														
	STQ	00010b														
Tag Error	Tag array	10000b				DRD, DWR, Probe	G	L1	1	0/1	1	-	-			
	STQ	10001b, 10011b				DWR			0/1	1	0/1					
	LDQ	10010b				DRD			0/1	1	0					
L1 TLB Error	TLB parity	00000b				TLB			-	D	L1	0/1	1	0	-	-
	TLB Multi-match	00001b										0/1	0			
	Locked TLB miss	00010b							1	1						
System Read Data Error	-	00000b	BUS	SRC	0	DRD	MEM/IO	LG	1	1	0	-	-			
Internal Error	Type 1	00001b		GEN	1	GEN	GEN	LG	1	0	0	-	-			
	Type 2	00010b														

MSR0000_0402 LS Machine Check Address (MC0_ADDR)

Read-write; updated-by-hardware. Cold reset: 0. The MCi_ADDR register contains valid data if indicated by MCi_STATUS[AddrV]. See 2.15.1 [Machine Check Architecture]. MSR0000_0000 is an alias of MSR0000_0402.

Bits	Description
63:0	ADDR: Address. See Table 214 .

Table 214: LS Address Register

Error Type	Error Sub-type	Address Register Bits	Description
System Read Data Error	-	47:6	Physical address[47:6]
Line Fill Error	-	47:6	Physical address[47:6]
Data Cache Error	Data array	47:4	Physical address[47:4]
	SCB	11:4	Physical address[11:4]
	STQ	4:0	Index

Table 214: LS Address Register

Error Type	Error Sub-type	Address Register Bits	Description
Tag Error	Tag array	47:4	Physical address[47:4] For Probe errors, [47:6]
	STQ	4:0	Index
	LDQ	5:0	Index
L1 TLB Error	TLB parity	47:12 4:0	Linear address[47:12] TLB index
	Locked TLB miss	47:12	Linear address[47:12]

MSR0000_0403 LS Machine Check Miscellaneous (MC0_MISC)

See 2.15.1.7 [Error Thresholding].

Bits	Description
63	Valid. IF (MSRC001_0015[McStatusWrEn]) THEN Read-write. ELSE Read-only. ENDIF. Reset: 1. 1=A valid CntP field is present in this register.
62	CntP: counter present. IF (MSRC001_0015[McStatusWrEn]) THEN Read-write. ELSE Read-only. ENDIF. Reset: 1. 1=A valid threshold counter is present.
61:52	Reserved.
51	CntEn: counter enable. Read-write. Reset: 0. 1=Count errors specified by the “Error Threshold Supported” column of Table 213.
50:49	Reserved.
48	Ovrflw: overflow. Read-write; set-by-hardware. Cold reset: 0. This bit is set by hardware when ErrCnt transitions from FFEh to FFFh.
47:44	Reserved.
43:32	ErrCnt: error counter. Read-write; updated-by-hardware. Cold reset: 0. This is written by software to set the starting value of the error counter. This is incremented by hardware when errors are logged. When this counter overflows, it stays at FFFh (no rollover). To set the threshold value, software should subtract the desired error count (the number of errors necessary in order for an interrupt to be taken) from FFFh and write the result into this field.
31:24	BlkPtr: Block pointer for additional MISC registers. Read-only. Reset: 00h. 00h=Extended MISC MSR block is not valid.
23:0	Reserved.

MSR0000_0404 IF Machine Check Control (MC1_CTL)

SharedC; Read-write. Reset: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture]. See MSRC001_0045 [IF Machine Check Control Mask (MC1_CTL_MASK)].

Bits	Description
63:24	Unused.
23	IVP: IC valid bit parity error.
22	L1TLBM: IC L1 TLB multi-match error.
21	L2TLBM: IC L2 TLB multi-match error.

Bits	Description
20	DFIFOE: decoder FIFO parity error.
19	DPDBE: decoder predecode buffer parity error.
18	DEIBP: decoder instruction buffer parity error.
17	DEUOPQP: Decoder micro-op queue parity error.
16	DEPRP: decoder patch RAM parity error.
15	BSRP: branch status register parity error.
14	Unused.
13	PQP: prediction queue parity error.
12	PFBP: prefetch buffer parity.
11:10	Unused.
9	SRDE: system read data error.
8	Unused.
7	LFE: line fill error. Uncorrectable error on cache line fill.
6	L1TP: L1 TLB parity error.
5	L2TP: L2 TLB parity error.
4	ISTP: L1 cache probe tag array parity error.
3	IMTP: L1 cache main tag array parity error.
2	IDP: L1 cache data array parity errors.
1:0	Unused.

MSR0000_0405 IF Machine Check Status (MC1_STATUS)

Cold reset: 0. See 2.15.1 [Machine Check Architecture]. See MSRC001_0015[McStatusWrEn]. Table 215 describes each error type. Table 216 describes the error codes and status register settings for each error type.

Bits	Description
63	Val: valid. See: MSR0000_0401[Val].
62	Overflow: error overflow. See: MSR0000_0401[Overflow].
61	UC: error uncorrected. See: MSR0000_0401[UC].
60	En: error enable. See: MSR0000_0401[En].
59	MiscV: miscellaneous error register valid. See: MSR0000_0401[MiscV]. 1=Valid thresholding in MSR0000_0407.
58	AddrV: error address valid. See: MSR0000_0401[AddrV].
57	PCC: processor context corrupt. See: MSR0000_0401[PCC].
56:45	Reserved.
44:43	Reserved.
42:40	Reserved.

39:36	Way: cache way in error. Read-write; Updated-by-hardware. Indicates the cache way in error.								
	<table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Way 0</td> </tr> <tr> <td>1h</td> <td>Way 1</td> </tr> <tr> <td>Fh-2h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	Way 0	1h	Way 1	Fh-2h	Reserved
<u>Bits</u>	<u>Description</u>								
0h	Way 0								
1h	Way 1								
Fh-2h	Reserved								
35:21	Reserved.								
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. See MSR0000_0401 [ErrorCodeExt] for description. See Table 216 for values.								
15:0	ErrorCode: error code. See: MSR0000_0401 [ErrorCode].								

Table 215: IF Error Descriptions

Error Type	Error Sub-type	Description	CTL ³	CID ²	EAC ¹
Line Fill Error	-	An uncorrectable error occurred during a demand line fill from the L2 cache; a machine check occurs before the instruction retires.	LineFill-Poison	A	E

Table 215: IF Error Descriptions

Error Type	Error Sub-type	Description	CTL ³	CID ²	EAC ¹
Instruction cache read error	IC Data Load Parity	A parity error occurred during load of data from the instruction cache. The data is discarded from the IC and can be refetched.	IDP	A	D
	IC valid bit	Parity error for IC valid bit.	IVP	A	D
	Main tag	A main tag parity error occurred.	IMTP	A	D
	Prediction queue	Parity error in prediction queue.	PQP	A	E
	PFB data/address	PFB data/address had a parity error. A PFB valid bit error, PFB multimatch error, Line Fill Error, or ReadData Error may additionally cause a PFB data/address error.	PFBP	A	E
	PFB valid bit	PFB valid bit had a parity error. This error may cause subsequent errors related to the entry, but the effect can be contained to the running process.		B	E
	PFB non-cacheable bit	PFB non-cacheable bit had a parity error.		B	E
	PFB promotion address error	An address parity error was detected when promoting from the PFB to the IC.		B	E
	Branch status register	A parity error was discovered in the branch status register. This error is uncorrectable, but the effect can be contained to the running process.	BSRP	A	E
Microcode Patch Buffer	Parity error in microcode patch buffer. This error is uncorrectable. If a reset is not performed or the patch area is not reloaded, then it is recommended that the compute unit be removed from the running configuration by the operating system if possible. After a reset, BIST is used to determine whether there is a hard fault in the buffer. If a hard fault is not found, the error was likely a transient upset and the buffer is not broken. This error can also be caused by an error in the microcode patch region of the CC6 save area.	DEPRP	A	E	
Instruction cache read error	Decoder micro-op queue	Parity error in decode unit. This error is correctable unless the operation is for a non-cacheable operand.	DEUQ	A	E
	Decoder instruction buffer		DEIBP	A	E
	Decoder pre-decode buffer		DEPD	A	E
	Decoder fetch address FIFO		DEFF	A	E

Table 215: IF Error Descriptions

Error Type	Error Sub-type	Description	CTL ³	CID ²	EAC ¹
Tag Probe	Probe tag error	A tag error was encountered during probe or victimization.	ISTP	0	D
	Probe tag valid bit	Parity error for IC probe tag valid bit.	IVP	0	D
L1 TLB	Parity	Parity error in L1 TLB.	L1TP	A	D
	Multimatch	Hit multiple entries in L1 TLB.	L1TLBM	A	D
L2 TLB	Parity	Parity error in L2 TLB.	L2TP	A	D
	Multimatch	Hit multiple entries in L2 TLB.	L2TLBM	A	D
System Read Data Error	-	An error occurred during an attempted demand read of data from the NB. Possible reasons include master abort, target abort.	SRDE	A	E
<p>1. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See 2.15.1.3.</p> <p>2. CID: core ID. A=Error reported to the affected core. 0=Error reported to core 0 of the compute unit. B=Error reported to both cores of the compute unit. See 2.15.1.3.</p> <p>3. See MSR0000_0404.</p>					

Table 216: IF Error Signatures

Error Type	Error Sub-type	Error-Code-Ext	Error Code						UC	ADDRV	PCC	Reserved	Reserved	
			Type	PP	T	RRRR	II/TT	LL						
Line Fill Error	-	00000b	MEM	-	-	IRD	I	L2	1	1	0	-	-	
Instruction Cache Read Error	IC data load parity	00001b				IRD	I	L1	0	1	0	-	-	
	IC valid bit	00010b							0	1	0	-	-	
	Main tag	00011b							0	1	0	-	-	
	Prediction queue	00100b							1	0	0	-	-	
	PFB data/address	00101b							0/1	0	0	-	-	
	PFB valid bit	01101b							1	0	0	-	-	
	PFB non-cacheable bit	01010b							0/1	0	0	-	-	
	PFB promotion address error	00111b							1	0	1	-	-	
	Branch status register	00110b							1	0	0	-	-	
	Micro-code Patch Buffer	10000b							LG	1	1	1	-	-
	Decoder micro-op queue	10001b							L1	0/1	1	0	-	-
	Decoder instruction buffer	10010b								0/1	1	0	-	-
	Decoder pre-decode buffer	10011b								0/1	0	0	-	-
	Decoder fetch address FIFO	10100b								0/1	1	0	-	-

Table 216: IF Error Signatures

Error Type	Error Sub-type	Error-Code-Ext	Error Code						UC	ADDRV	PCC	Reserved	Reserved
			Type	PP	T	RRRR	II/TT	LL					
Tag Probe	Probe tag error	01000b	MEM	-	-	Probe	I	L1	0	1	0	-	-
	Probe tag valid bit	01001b											
L1 TLB	Parity	00000b	TLB	-	-	-	I	L1	0	1	0	-	-
	Multi-match	00001b											
L2 TLB	Parity	00000b						L2					
	Multi-match	00001b											
System Read Data Error	-	00000b	BUS	SRC	0	IRD	MEM	LG	1	1	0	-	-

MSR0000_0406 IF Machine Check Address (MC1_ADDR)

Read-write; updated-by-hardware. Cold reset: 0. The MCi_ADDR register contains valid data if indicated by MCi_STATUS[AddrV]. See 2.15.1 [Machine Check Architecture].

Bits	Description
63:0	ADDR. See Table 217.

Table 217: IF Address Register

Error Type	Error Sub-Type	Address Register Bits	Description
Line Fill Error	-	[63:48] [47:6] [5:0]	Reserved Linear address. Reserved
Patch RAM parity	-	[63:4] [3:0]	Reserved Line group index.

Table 217: IF Address Register

Error Type	Error Sub-Type	Address Register Bits	Description
Instruction cache read error	IC data load parity	[63:48] [47:6]	Reserved Linear address
	IC valid bit	[7:6]	Bank
	Main tag	[1]	Way
		[5:2,0]	Reserved
	Decoder micro-op queue parity	[63:4] [3:2] [1:0]	Reserved Micro-op queue in error bitmask: <ul style="list-style-type: none"> • [3]: Dispatch micro-op queue • [2]: FDC micro-op queue Micro-op queue slot in error.
Decoder instruction buffer parity	[63:2] [1:0]	Reserved Bank and parity bit in error: <ul style="list-style-type: none"> • 00b: Bank A, parity bit 0 or 1 • 01b: Bank B, parity bit 0 or 1 • 10b: Bank A, parity bit 2 or 3 • 11b: Bank B, parity bit 2 or 3 	
Decoder FIFO parity	[63:1] [0]	Reserved Bank in error: 0=Bank A. 1=Bank B.	
Tag Probe	-	[63:48] [47:6] [5:4] [3:0]	Reserved Physical address. Reserved Bank bitmask: <ul style="list-style-type: none"> • [0]: Bank 0 • [1]: Bank 1 • [2]: Bank 2 • [3]: Bank 3
L1 TLB	-	[63:48] [47:12]	Reserved Linear address.
L1 TLB Multi-match	-	[11:3] [2:0]	4-KB page: <ul style="list-style-type: none"> • [47:12]: Linear address. 2-MB page: <ul style="list-style-type: none"> • [47:20]: Linear address. • [19:12]: Reserved Reserved Bank bitmask: <ul style="list-style-type: none"> • [0]: Bank 0 • [1]: Bank 1 • [2]: Bank 2
L2 TLB	-	[63:48]	Reserved
L2 TLB Multi-match	-	[47:12] [11:4] [3:0]	Linear address. (4-KB page size only) Reserved Match lines.

MSR0000_0407 IF Machine Check Miscellaneous (MC1_MISC)

See [2.15.1.7 \[Error Thresholding\]](#).

Bits	Description
63	Valid. See: MSR0000_0403[Valid] .
62	CntP: counter present. See: MSR0000_0403[CntP] .
61:52	Reserved.
51	CntEn: counter enable. See: MSR0000_0403[CntEn] .
50:49	Reserved.
48	Ovrflw: overflow. See: MSR0000_0403[Ovrflw] .
47:44	Reserved.
43:32	ErrCnt: error counter. See: MSR0000_0403[ErrCnt] .
31:24	BlkPtr: Block pointer for additional MISC registers. See: MSR0000_0403[BlkPtr] .
23:0	Reserved.

MSR0000_0408 CU Machine Check Control (MC2_CTL)

SharedC; Read-write. Reset: 0000_0000_0000_0000h. See [2.15.1 \[Machine Check Architecture\]](#). See [MSRC001_0046 \[CU Machine Check Control Mask \(MC2_CTL_MASK\)\]](#).

Bits	Description
63:15	Unused.
14	L2TlbFill: TLB fill error enable. Data with uncorrectable error provided to TLB.
13	RdData: read data error from NB.
12	L2Tag: L2 cache tag error.
11	L2TlbData: L2 TLB parity error. Parity error reading from TLB.
10	L2Prefetch: L2 data prefetcher parity error.
9	XabAddr: XAB address parity error.
8	PrbAddr: probe buffer address parity error.
7	FillData: fill data parity and ECC error.
6	PrqAddr: post retire queue address parity error.
5	PrqData: post retire queue data parity error.
4	WccAddr: write coalescing cache address ECC error.
3	WccData: write coalescing cache data ECC error.
2	WcbData: write combining buffer data parity error.
1	VbData: victim buffer data parity and ECC error.
0	L2TagMultiHit: L2 tag multiple hit error.

MSR0000_0409 CU Machine Check Status (MC2_STATUS)

Cold reset: 0. See [2.15.1 \[Machine Check Architecture\]](#). See [MSRC001_0015\[McStatusWrEn\]](#). [Table 219](#) describes each error type. [Table 220](#) describes the error codes and status register settings for each error type.

Bits	Description										
63	Val: valid. See: MSR0000_0401[Val] .										
62	Overflow: error overflow. See: MSR0000_0401[Overflow] .										
61	UC: error uncorrected. See: MSR0000_0401[UC] .										
60	En: error enable. See: MSR0000_0401[En] .										
59	MiscV: miscellaneous error register valid. See: MSR0000_0401[MiscV] . 1=Valid thresholding in MSR0000_040B .										
58	AddrV: error address valid. See: MSR0000_0401[AddrV] .										
57	PCC: processor context corrupt. See: MSR0000_0401[PCC] .										
56:55	Reserved.										
54:47	Syndrome[7:0]. Read-write. Syndrome[11:0] = {Syndrome[11:8], Syndrome[7:0]}. The syndrome bits when an ECC error is detected. See Table 220 for when Syndrome[11:0] is valid. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Array</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>L2 Tag</td> <td>Syndrome[7:0].</td> </tr> <tr> <td>WCC Tag</td> <td>Syndrome[11:0].</td> </tr> <tr> <td>L2 Data</td> <td>Syndrome[8:0].</td> </tr> <tr> <td>WCC Data</td> <td>Syndrome[8:0].</td> </tr> </tbody> </table>	Array	Description	L2 Tag	Syndrome[7:0].	WCC Tag	Syndrome[11:0].	L2 Data	Syndrome[8:0].	WCC Data	Syndrome[8:0].
Array	Description										
L2 Tag	Syndrome[7:0].										
WCC Tag	Syndrome[11:0].										
L2 Data	Syndrome[8:0].										
WCC Data	Syndrome[8:0].										
46	CECC: correctable ECC error. Read-write; Updated-by-hardware. 1=The error was a correctable ECC error according to the restrictions of the ECC algorithm. UC indicates whether the error was actually corrected by the processor.										
45	UECC: uncorrectable ECC error. Read-write; Updated-by-hardware. 1=The error was an uncorrectable ECC error according to the restrictions of the ECC algorithm. UC indicates whether the error was actually corrected by the processor.										
44:43	Reserved.										
42:40	Reserved.										
39:36	Way: cache way in error. Read-write; Updated-by-hardware. Indicates the cache way in error. See Table 220 for when Way is valid and what ways are valid. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Way 0</td> </tr> <tr> <td>1h</td> <td>Way 1</td> </tr> <tr> <td>Eh-2h</td> <td>Way <Way></td> </tr> <tr> <td>Fh</td> <td>Way 15</td> </tr> </tbody> </table>	Bits	Description	0h	Way 0	1h	Way 1	Eh-2h	Way <Way>	Fh	Way 15
Bits	Description										
0h	Way 0										
1h	Way 1										
Eh-2h	Way <Way>										
Fh	Way 15										
35:28	Reserved.										
27:24	Syndrome[11:8]. See: MSR0000_0409[Syndrome[7:0]] .										
23:21	Reserved.										
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. See MSR0000_0401[ErrorCodeExt] for description. See Table 220 for values.										
15:0	ErrorCode: error code. See: MSR0000_0401[ErrorCode] .										

Table 218: MB, SBU, and SBC Definitions

Term	Definition
MB	Multi-bit ECC error, uncorrected.
SBU	Single-bit ECC error, not-corrected. There are some implementation specific conditions when a single bit error is not correctable.
SBC	Single-bit ECC error is detected and correctable.

Table 219: CU Error Descriptions

Error Type	Error Sub-Type	Description	CTL ³	CID ²	EAC ¹
System Read Data	TLB	An error occurred during an attempted read of data from the NB. Possible reasons include master abort, target abort, and receipt of read data error for TLB. Error Action ⁴ : None.	RdData	A	D
	Prefetch				
	Wcc				
TLB	TlbPar	Data parity error reading from TLB. Error Action ⁴ : Invalidate TLB entry.	L2TlbData	A	D
	FillErr	Poison data provided for TLB fill. Error Action ⁴ : None.	L2TlbPoison	A	D

Table 219: CU Error Descriptions

Error Type	Error Sub-Type	Description	CTL ³	CID ²	EAC ¹
L2 Cache	Prefetch	Prefetcher request FIFO parity error. Error Action ⁴ : Invalidate entry (drop prefetch).	L2Prefetch	A	D
	FillEcc	Fill ECC error on data fills. <u>SubCase</u> <u>Error Action⁴</u> MB, SBU Poison data returned to destination; error remains in source. SBC Corrected data returned to destination; error remains in source. The data sources are indicated in LL field and affect what part of Way is valid: <u>Source</u> <u>LL</u> <u>Way</u> WCC L1 [1:0]See Note 1. L2 L2 [3:0] NB LG - See Note 2. Notes: 1. WCC: Indicates data corrupted in WCC or Fill Buffer. 2. NB: Indicates data corrupted in Fill Buffer.	FillData	A	D
	FillPar	Fill parity error on instruction fills. <u>SubCase</u> <u>Description</u> NB->IC NB to IC parity error: Error Action ⁴ : Invalidate data and Nack request (IC will re-request). L2->IC L2 to IC parity error: Error Action ⁴ : Invalidate data and Nack request (IC will re-request). L2->LS,TLB L2 to LS or TLB parity error: Error Action ⁴ : Poison data returned to destination; error remains in source.	FillData	A	D
	PrqAddr	Post Retire Queue address parity error. Error Action ⁴ : Sync flood.	PrqAddr	A	D
	PrqData	Post Retire Queue data parity error. Error Action ⁴ : Poison line WCC or line sent to NB.	PrqData	A	D
	L2 Cache	WccTag	Write Coalescing Cache tag ECC error. <u>SubCase</u> <u>Error Action⁴</u> MB, SBU Sync flood. SBC Invalidate Wcc tag entry (cleans error).	WccAddr	0
WccData		WCC data ECC error. <u>SubCase</u> <u>Error Action⁴</u> MB, SBU Poison copy in WCC. SBC Corrected copy in WCC.	WccData	A	D
WcbData		WCB data parity error. Error Action ⁴ : Poison sent to NB.	WcbData	A	D

Table 219: CU Error Descriptions

Error Type	Error Sub-Type	Description	CTL ³	CID ²	EAC ¹
L2 Cache (continued)	VbData	<p>VB data ECC or parity error.</p> <p><u>SubCase</u> <u>Description</u></p> <p>Par Parity: Parity error indicated when CECC and UECC are both clear. Error Action⁴: Poison sent to NB.</p> <p>MB,SBU Error Action⁴: Poison sent to NB.</p> <p>SBC Single-bit ECC error, corrected: Error Action⁴: Corrected data sent to NB.</p>	VbData	0	D
Tag	L2TagMH	Multiple hits on L2 tag. Error Action ⁴ : Sync flood.	L2TagMulti Hit	0	D
	L2Tag	<p>A correctable or uncorrectable ECC error was seen in the L2 tag. The L2TagMH error signature supersedes the L2Tag error signature if they both occur for the same L2 tag read.</p> <p><u>SubCase</u> <u>Description</u></p> <p>MB, SBU Error Action⁴: Sync flood.</p> <p>SBC Error Action⁴: Correct error in array and retry the operation.</p> <p>Hard A hard error was seen in the L2 tag. Error Action⁴: Sync flood.</p>	L2Tag	0	D
	XabAddr	Transaction Address Buffer (XAB) parity error. This error is system fatal; memory coherence may have been affected. Error Action ⁴ : Sync flood.	XabAddr	A	D
	PrbAddr	Probe buffer address parity error. This error is system fatal; memory coherence may have been affected. Error Action ⁴ : Sync flood.	PrbAddr	0	D

1. EAC: The error action is taken if detected for all CU errors. D=Error action taken if detected. E=Error action taken if MCA bank enabled. See 2.15.1.3.
2. CID: core ID. A=Error reported to the affected core. 0=Error reported to core 0 of the compute unit; see 2.15.1.3.
3. See MSR0000_0408.
4. Error Action: Sync flood=Take sync flood if PCC=1. None=No action other than that specified by MCA.

Table 220: CU Error Signatures

Error Type	Error Sub-Type	Sub Case	Error-Code-Ext	Error Code						UC	ADDR	PCC	Syn-drome	Way	CECC	UECC	Deferred	Poison
				Type	PP	T	RRRR	II/TT	LL									
TLB	TlbPar	-	00000b	TLB	-	-	-	G	L2	0	1	0	-	[2:0]	0	0	0	0
	FillErr	-	00001b							1				-				1
System Read Data	TLB	-	00000b	BUS	SRC	0	RD	MEM/IO	L2	1	1	0	-	-	0	0	0	0
	Prefetch		00001b							0								
	Wcc		00010b				DWR	MEM	L1	1								

Table 220: CU Error Signatures

Error Type	Error Sub-Type	Sub Case	Error-Code-Ext	Error Code						UC	ADDR	PCC	Syn-drome	Way	CECC	UECC	Deferred	Poison													
				Type	PP	T	RRRR	I/TT	LL																						
L2 Cache	FillEcc	MB	00100b	MEM	-	-	DRD	D	See ¹	0	1	0	[8:0]	See ¹	0	1	1	0													
		SBU													1	0															
		SBC															0														
L2 Cache	FillPar	NB->IC	00101b	MEM	-	-	IRD	I	LG	0	0	0	-	-	0	0	0	0													
		L2->IC																	D	L2	0	1	0	[3:0]	1	0					
		L2->LS,TLB																													
L2 Cache	Prefetch	-	00110b	MEM	-	-	Prefetch	D	L2	0	1	0	-	-	0	0	0	0													
	PrqAddr	-																	00111b	DWR	D	L1	1	0	1	-	-	0	0	0	0
	PrqData	-																													
L2 Cache	WccTag	MB	01001b	MEM	-	-	DWR	D	L1	1	1	1	[11:0]	[1:0]	0	1	0	0													
		SBU													1	0															
		SBC													0		0														
	WccData	MB	01010b	0	0	[8:0]	0	1	1	0																					
		SBU					1	0	1	0																					
		SBC							0																						
WcbData	-	01011b			LG	0		0	-	-	0	0	1																		
L2 Cache	VbData	Par	01100b	MEM	-	-	Probe, Evict	I	L2	0	0	0	-	-	0	0	1	0													
		MB													[8:0]	0	1														
		SBU														1	0														
		SBC																0													
Tag	L2Tag	MB	10000b	MEM	-	-	GEN	G	L2	1	1	1	[7:0]	[3:0]	0	1	0	0													
		SBU								1		1																			
		SBC								0		0																			
		Hard								10001b	1				1	-	0	0													
	L2Tag MH	-	10010b			1	1	1	-	-	0	0	0	0																	
	XabAddr	-	10011b						-	-																					
	PrbAddr	-	10100b						Probe				-	-																	

1. LL and Way are specified in Table 219.

MSR0000_040A CU Machine Check Address (MC2_ADDR)

Read-write; updated-by-hardware. Cold reset: 0. The MCi_ADDR register contains valid data if indicated by MCi_STATUS[AddrV]. See 2.15.1 [Machine Check Architecture].

Bits	Description
63:0	ADDR. See Table 221.

Table 221: CU Address Register

Error Type	Error Sub-Type	SubCase	Description	
System Read Data Error	-	-	<u>Bit</u>	<u>Description</u>
			[63:48]	Reserved.
			[47:6]	PhysAddr[47:6]
L2 Cache	TlbPar	-	<u>Bit</u>	<u>Description</u>
			[63:7]	Reserved.
	FillErr	-	[6:0]	Index[6:0]
			<u>Bit</u>	<u>Description</u>
			[63:48]	Reserved.
			[47:6]	PhysAddr[47:6]
L2 Cache	Prefetch	-	<u>Bit</u>	<u>Description</u>
			[63:5]	Reserved.
	FillEcc	All	[4:0]	Prefetch FIFO read pointer.
			<u>Bit</u>	<u>Description</u>
			[63:48]	Reserved.
			[47:3]	PhysAddr[47:3].
			FillPar	L2->LS,TLB
[63:48]	Reserved.			
WccTag	All	[47:3]	PhysAddr[47:3].	
		[2:0]	Reserved.	
		<u>Bit</u>	<u>Description</u>	
WccData	All	[63:10]	Reserved.	
		[9:6]	Index[9:6]	
WcbData	-	[5:0]	Reserved.	
		<u>Bit</u>	<u>Description</u>	
		[63:48]	Reserved.	
		[47:3]	PhysAddr[47:3]	
WcbData	-	[2]	Reserved.	
		[1:0]	Index[1:0]	

Table 221: CU Address Register

Error Type	Error Sub-Type	SubCase	Description	
Tag	L2TagMH	-	<u>Bits</u>	<u>Description</u>
	L2Tag	All	[63:17]	Reserved.
			[16:6]	PhysAddr[16:6]. (dependent on cache size ¹) [16:6]: 2 MB [15:6]: 1 MB ([16] Reserved)
			[5:0]	Reserved.
	XabAddr	-	<u>Bits</u>	<u>Description</u>
			[63:5]	Reserved.
			[4:0]	XAB index.
	PrbAddr	-	<u>Bits</u>	<u>Description</u>
			[63:4]	Reserved.
			[3:0]	Probe buffer index.

1. See [CPUID Fn8000_0006_ECX\[L2Size\]](#).

MSR0000_040B CU Machine Check Miscellaneous (MC2_MISC)See [2.15.1.7 \[Error Thresholding\]](#).

Bits	Description
63	Valid. See: MSR0000_0403[Valid] .
62	CntP: counter present. See: MSR0000_0403[CntP] .
61:52	Reserved.
51	CntEn: counter enable. See: MSR0000_0403[CntEn] .
50:49	Reserved.
48	Ovrflw: overflow. See: MSR0000_0403[Ovrflw] .
47:44	Reserved.
43:32	ErrCnt: error counter. See: MSR0000_0403[ErrCnt] .
31:24	BlkPtr: Block pointer for additional MISC registers. See: MSR0000_0403[BlkPtr] .
23:0	Reserved.

MSR0000_040C MC3 Machine Check Control (MC3_CTL)

Reset: 0000_0000_0000_0000h. Read-only.

Bits	Description
63:0	Unused.

MSR0000_040D MC3 Machine Check Status (MC3_STATUS)Reset: 0. See [MSRC001_0015\[McStatusWrEn\]](#).

Bits	Description
63:0	Reserved.

MSR0000_040E MC3 Machine Check Address (MC3_ADDR)

Reset: 0000_0000_0000_0000h. Read-only.

Bits	Description
63:0	Reserved.

MSR0000_040F MC3 Machine Check Miscellaneous (MC3_MISC)

Reset: 0000_0000_0000_0000h. Read-only.

Bits	Description
63:0	Reserved.

MSR0000_0410 NB Machine Check Control (MC4_CTL)

Read-write; **Not-same-for-all**. Reset: 0000_0000_0000_0000h. MSR0000_0410[31:0] is an alias of D18F3x40. This register is also accessible through PCI configuration space; see D18F3x40 [MCA NB Control]. Only one of these registers exists in multi-core devices; see 3.1.1 [Northbridge MSRs In Multi-Core Products]. See D18F3x44 [MCA NB Configuration]. See 2.15.1 [Machine Check Architecture]. Reporting is also masked by MSRC001_0048 [NB Machine Check Control Mask (MC4_CTL_MASK)]. Accessibility of this register by non-NBC cores is affected by D18F3x44[NbMcaToMstCpuEn].

Bits	Description
63:32	Unused.
31	McaCpuDatErrEn: compute unit data error. 1=Enables MCA reporting of CPU data errors sent to the NB.
30	Unused.
29:27	Unused.
26	NbArrayParEn: northbridge array parity error reporting enable. 1=Enables reporting of parity errors in the NB arrays.
25	UsPwDatErrEn: upstream data error enable. Read-write. 1=Enables MCA reporting of upstream posted writes in which the EP bit is set.
24:18	Unused.
17	CpPktDatEn: completion packet error reporting enable. Read-write. 1=Enables MCA reporting of completion packets with the EP bit set.
16	NbIntProtEn: northbridge internal bus protocol error reporting enable. Read-write. 1=Enables MCA reporting of protocol errors detected on the northbridge internal bus. When possible, this enable should be cleared before initiating a warm reset to avoid logging spurious errors due to RESET# signal skew.
15:13	Unused.

12	WDTRptEn: watchdog timer error reporting enable. 1=Enables MCA reporting of watchdog timer errors. The watchdog timer checks for NB system accesses for which a response is expected but no response is received. See D18F3x44 [MCA NB Configuration] for information regarding configuration of the watchdog timer duration. This bit does not affect operation of the watchdog timer in terms of its ability to complete an access that would otherwise cause a system hang. This bit only affects whether such errors are reported through MCA.
11	AtomicRMWEn: atomic read-modify-write error reporting enable. 1=Enables MCA reporting of atomic read-modify-write (RMW) commands received from an IO link. Atomic RMW commands are not supported. An atomic RMW command results in a link error response being generated back to the requesting IO device. The generation of the link error response is not affected by this bit.
10	Unused.
9	TgtAbortEn: target abort error reporting enable. 1=Enables MCA reporting of target aborts to a link. The NB returns an error response back to the requestor with any associated data all 1s independent of the state of this bit.
8	MstrAbortEn: master abort error reporting enable. 1=Enables MCA reporting of master aborts to a link. The NB returns an error response back to the requestor with any associated data all 1s independent of the state of this bit.
7:6	Unused.
5	SyncPktEn: link sync packet error reporting enable. 1=Enables MCA reporting of link-defined sync error packets detected on link. The NB floods its outgoing link with sync packets after detecting a sync packet on the incoming link independent of the state of this bit.
4:0	Unused.

MSR0000_0411 NB Machine Check Status (MC4_STATUS)

Not-same-for-all. Cold reset: 0000_0000_0000_0000h. See [2.15.1 \[Machine Check Architecture\]](#). Only one of these registers exists in multi-core devices; see [3.1.1 \[Northbridge MSRs In Multi-Core Products\]](#). See [MSRC001_0015\[McStatusWrEn\]](#). [Table 222](#) describes each error type. [Table 223](#) and [Table 224](#) describe the error codes and status register settings for each error type. [MSR0000_0411\[31:0\]](#) is an alias of [D18F3x48](#). [MSR0000_0411\[63:32\]](#) is an alias of [D18F3x4C](#). Accessibility of this register by non-NBC cores is affected by [D18F3x44\[NbMcaToMstCpuEn\]](#).

Bits	Description
63	Val: valid. See: MSR0000_0401[Val] .
62	Overflow: error overflow. See: MSR0000_0401[Overflow] .
61	UC: error uncorrected. See: MSR0000_0401[UC] .
60	En: error enable. See: MSR0000_0401[En] .
59	MiscV: miscellaneous error register valid. See: MSR0000_0401[MiscV] . 1=Valid thresholding in MSR0000_0413 or MSRC000_0408 .
58	AddrV: error address valid. See: MSR0000_0401[AddrV] .
57	PCC: processor context corrupt. See: MSR0000_0401[PCC] .
56	ErrCoreIdVal: error core ID is valid. Read-write; set-by-hardware. 1=The ErrCoreId field is valid.
55	Reserved.
54:45	Reserved.
44	Reserved.

43:42	Reserved.
41	SubLink: sublink. Read-write; set-by-hardware. For errors associated with a link, this bit indicates if the error was associated with the upper or lower byte of the link. 0=Sublink [7:0]. 1=Sublink [15:8].
40	Reserved.
39:37	Reserved.
36	Link. Read-write; set-by-hardware. For errors associated with a link, this field indicates that the link was associated with the error.
35:32	ErrCoreId: error associated with core N. Read-write; updated-by-hardware. When ErrCoreIdVal=1 this field indicates which core within the processor is associated with the error; Otherwise this field is reserved. All values greater than D18F5x84[CmpCap] are reserved.
31:24	Reserved.
23:21	Reserved.
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. See MSR0000_0401[ErrorCodeExt] for description. See Table 223 for values.
15:0	ErrorCode: error code. See: MSR0000_0401[ErrorCode].

Table 222: NB Error Descriptions

Error Type	Description	CTL ¹	ETG ²	EAC
Sync Error	Link-defined sync error packets detected on link. The NB floods its outgoing links with sync packets after detecting a sync packet on an incoming link independent of the state of the control bits.	SyncPktEn	L	D
Master Abort	Master abort seen as result of link operation. Reasons for this error include requests to non-existent addresses. The NB returns an error response back to the requestor with any associated data all 1s independent of the state of the control bit.	MstrAbortEn	L	D
Target Abort	Target abort seen as result of link operation. The NB returns an error response back to the requestor with any associated data all 1s independent of the state of the control bit.	TgtAbortEn	L	D
RMW Error	An atomic read-modify-write (RMW) command was received from an IO link. Atomic RMW commands are not supported. An atomic RMW command results in a link error response being generated back to the requesting IO device. The generation of the link error response is not affected by the control bit.	AtomicRMWEn	L	D
WDT Error	NB WDT timeout due to lack of progress. The NB WDT monitors transaction completions. A transaction that exceeds the programmed time limit reports errors via the MCA. The cause of error may be another node or device which failed to respond.	WDTRptEn	L	D
Link Data Error	Data error detected on link. If enabled for reporting and the request is sourced from a core, then PCC is set. (If not enabled for reporting, PCC is not set. If configured to allow an error response to be returned to the core, this could allow error containment to a scope smaller than the entire system.)	McaUsPwDat-ErrEn, CpPkt-DatEn	L	D

Table 222: NB Error Descriptions

Error Type	Description	CTL ¹	ETG ²	EAC
Protocol Error	<p>Protocol error detected by link. These errors are distinguished from each other by the value in MSR0000_0412[ErrAddr]. See Table 226.</p> <p>For protocol errors, the system cannot continue operation. Protocol errors can be caused by other subcomponents than the one reporting the error. For diagnosis, collect and examine MCA registers from other banks, cores, and processors in the system.</p>	NbIntProtEn	L ³	D
NB Array Error	A parity error was detected in the NB internal arrays.	NbArrayParEn	-	D
Compute Unit Data Error	<p>NB received a data error from a core and this error could not be contained. Sync flood will occur if D18F3x180[CpuLeakErrEn] is set. For the cause of the data error, examine the core MCA registers for deferred errors. This error may occur for the following types of data writes:</p> <ul style="list-style-type: none"> • DRAM • APIC • Configuration space (IO and MMIO) <p>For these errors, sync flood will occur if D18F3x180[CpuLeakErrEn] is set.</p>	McaCpuDat-ErrEn	-	D

1. See [MSR0000_0410](#).
2. ETG: error threshold group. L=Link. D=DRAM. C=L3 Cache. See [2.15.1.7 \[Error Thresholding\]](#).
3. The error thresholding group is Link if link protocol error; none for non-link protocol error.
4. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See [2.15.1.3](#).

The NB is capable of reporting the following errors:

Table 223: NB Error Signatures, Part 1

Error Type	Error Code-Ext	Error Code					
		Type	PP	T	RRRR	II/TT	LL
Reserved	00000b	-	-	-	-	-	-
Reserved	00001b	-	-	-	-	-	-
Sync Error	00010b	BUS	OBS	0	GEN	GEN	LG
Mst Abort	00011b	BUS	SRC/OBS	0	RD/WR	MEM/IO ¹	LG
Tgt Abort	00100b	BUS	SRC/OBS	0	RD/WR	MEM/IO ¹	LG
RMW Error	00110b	BUS	OBS	0	GEN	IO	LG
WDT Error	00111b	BUS	GEN	1	GEN	GEN	LG
Link Data Error	01010b	BUS	SRC/OBS	0	RD/WR/DWR	MEM/IO	LG
NB Protocol Error	01011b	BUS	OBS	0	GEN	GEN	LG
NB Array Error	01100b	BUS	OBS	0	GEN	GEN	LG
Compute Unit Data Error	11001b	MEM	-	-	WR	Data	LG

1. Indicates the type of link attached to the reporting NB, not the instruction type. MEM indicates coherent link, IO indicates IO link.

Table 224: NB Error Signatures, Part 2

Error Type	UC	AddrV	PCC	Syndrome Valid	CECC	UECC	Reserved	Scrub	Link	Err CoreId
Sync Error	1	0	1	-	0	0	-	0	Y	-
Mst Abort	1	1	Core ¹⁰	-	0	0	-	0	Y	Y
Tgt Abort	1	1	Core ¹⁰	-	0	0	-	0	Y	Y
RMW Error	1	1	0	-	0	0	-	0	Y	-
WDT Error	1	0 ¹	1	-	0	0	-	0	-	-
Link Data Error	1	1	0	-	0	0	-	0	Y	-
NB Protocol Error	1	1/0 ²	1	-	0	0	-	0	Y	-
NB Array Error	1	1 ⁴	1	-	0	0	-	0	-	-
Compute Unit Data Error	1	0	1	-	0	0	-	0	-	Y

1. See Table 229: [NB Address Register for Watchdog Timer Errors]
2. See Table 226 [NB Address Register for Protocol Errors]
3. Link identified only if link protocol error. See entry in Table 222.
4. See Table 228: [NB Address Register for NB Array Errors]
5. Depends on Memory Transaction Type (Table 61); valid if non-zero.
6. MS: multi-symbol. 1=Multi-symbol. 0=Not multi-symbol.
7. Retries initiated by either side of the link are logged.
8. MB: multi-bit error. 1=Multi-bit error.
9. SB: single bit error. 1=Single bit error.
10. Core: source is core. 1=Source is core. 0=Source is not core.
11. Deferred: error is deferred. 1=Error is deferred. 0=Error is not deferred.

MSR0000_0412 NB Machine Check Address (MC4_ADDR)

Read-write; Not-same-for-all. Cold reset: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture]. MSR0000_0412[31:0] is an alias of D18F3x50. MSR0000_0412[63:32] is an alias of D18F3x54. Only one of these registers exists in multi-core devices; see 3.1.1 [Northbridge MSRs In Multi-Core Products]. ErrAddr[47:1] carries supplemental information associated with a machine check error, generally the address being accessed. Accessibility of this register by non-NBC cores is affected by D18F3x44[NbMcaToMst-CpuEn].

Bits	Description
63:48	Reserved.
47:32	ErrAddr[47:32]: Error Address Bits[47:32]. See: ErrAddr[31:1].
31:1	ErrAddr[31:1]: Error Address Bits[31:1]. ErrAddr[63:0] = {00h, MSR0000_0412[ErrAddr[47:32]], MSR0000_0412[ErrAddr[31:1]], 0b}. See the tables below for the encoding.
0	Reserved.

The register format depends on the type of error being logged:

- Protocol errors contain the error reason code, may contain the physical address, and are formatted according to [Table 226](#).
- NB array errors indicate the array in error, and are formatted according to [Table 228](#).
- NB Watchdog timer errors depend on the mode selected by [D18F3x180\[McaLogErrAddrWdtErr\]](#), and the format is indicated by [D18F3x4C\[AddrV\]](#). If [D18F3x4C\[AddrV\]](#) is indicated, errors are formatted according to [Table 225](#). If [D18F3x4C\[AddrV\]](#) is not indicated, errors are formatted according to [Table 229](#).
- All other NB errors which indicate [D18F3x4C\[AddrV\]](#) are formatted according to [Table 225](#).

Table 225: NB Address Register Default Encoding

ErrAddr	Description
47:1	Physical address bits 47:1.

Table 226: NB Address Register for Protocol Errors

ErrAddr	Description
47:6	If D18F3x4C[AddrV] set, contains physical address bits 47:6, else reserved.
5:1	ProtocolErrorType. See Table 227

Table 227: Protocol Error Type

Link Protocol Error Type	Description
00000b	Link: SRQ Read Response without matching request
00001b	Link: Probe Response without matching request
00010b	Link: TgtDone without matching request
00011b	Link: TgtStart without matching request
00100b	Link: Command buffer overflow
00101b	Link: Data buffer overflow
00110b	Link: Link retry packet count acknowledge overflow
00111b	Link: Data command in the middle of a data transfer
01000b	Link: Link address extension command followed by a packet other than a command with address.
01001b	Link: A specific coherent-only packet from a CPU was issued to an IO link.
01010b	Link: A command with invalid encoding was received. This error occurs when: <ol style="list-style-type: none"> 1. Any invalid command is received (including a command with no valid encoding or a coherent link command over an IO link or vice versa) while not in retry mode. 2. Any illegal command is received in which the CRC is correct while in retry mode (including any upstream broadcast command (HT command encoding = 6'b11101x)).
01011b	Link: Link CTL deassertion occurred when a data phase was not pending. This error condition may only occur when error-retry mode is not enabled (if it is enabled, this condition triggers a retry).
011xxb	Reserved
10000b - 11111b	Reserved

Table 228: NB Address Register for NB Array Errors

ErrAddr bits	Array Code	Description
47:6	-	Reserved
5:1	00000b	SRA: System request address
	00001b	SRD: System request data
	00010b	SPB: System packet buffer
	00011b	MCD: Memory controller data
	00100b	MPB: Memory packet buffer
	00101b	LPB0: Link 0 packet buffer
	00110b-01000b	Reserved.
	01001b	MPBC: Memory controller command packet buffer
	01010b	MCDBM: Memory controller byte mask
	01011b	MCACAM: Memory controller address array
	01100b	DMAP: Extended DRAM address map
	01101b	MMAP: Extended MMIO address map
	01110b	X86MAP: Extended PCI/IO address map
	01111b	CFGMAP: Extended config address map
	10000b-10111b	Reserved.
	11000b	SRIMCTRTE: SRI/MCT extended routing table
	11001b-11100b	Reserved
	11101b	TCB: TCB array.
	11110b-11111b	Reserved

Table 229: NB Address Register for Watchdog Timer Errors

Bits	Description						
63:48	Reserved						
47:40	CoreId. Indicates the core ID if the SourcePointer specifies Core. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>07h-00h</td> <td>CoreId</td> </tr> <tr> <td>FFh-08h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	07h-00h	CoreId	FFh-08h	Reserved
<u>Bits</u>	<u>Description</u>						
07h-00h	CoreId						
FFh-08h	Reserved						
39:36	SystemResponseCount. This field records unspecified, implementation-specific information.						
35:31	WaitCode. records unspecified, implementation-specific information (all zeroes means no waiting condition).						
30	WaitForPostedWrite.						
29:27	DestinationNode. Records the Node ID of the node addressed by the transaction.						

Table 229: NB Address Register for Watchdog Timer Errors

Bits	Description																		
26:25	DestinationUnit. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Core</td> </tr> <tr> <td>01b</td> <td>Reserved</td> </tr> <tr> <td>10b</td> <td>Memory Controller</td> </tr> <tr> <td>11b</td> <td>Host</td> </tr> </tbody> </table>	Bits	Description	00b	Core	01b	Reserved	10b	Memory Controller	11b	Host								
Bits	Description																		
00b	Core																		
01b	Reserved																		
10b	Memory Controller																		
11b	Host																		
24:22	SourceNode. Records the Node ID of the node originating the transaction.																		
21:20	SourceUnit. (same encoding as Destination Unit)																		
19:15	SourcePointer. Identifies crossbar source: <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00000b</td> <td>SRI HostBridge</td> </tr> <tr> <td>00011b-00001b</td> <td>Reserved</td> </tr> <tr> <td>00100b</td> <td>Core. See CoreId.</td> </tr> <tr> <td>00111b-00101b</td> <td>Reserved</td> </tr> <tr> <td>01000b</td> <td>Memory controller.</td> </tr> <tr> <td>01111b-01001b</td> <td>Reserved</td> </tr> <tr> <td>1HH0Nb</td> <td>Link. Link HH; sublink N (where N=0b for ganged links)</td> </tr> <tr> <td>1xx1xb</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00000b	SRI HostBridge	00011b-00001b	Reserved	00100b	Core. See CoreId.	00111b-00101b	Reserved	01000b	Memory controller.	01111b-01001b	Reserved	1HH0Nb	Link. Link HH; sublink N (where N=0b for ganged links)	1xx1xb	Reserved
Bits	Description																		
00000b	SRI HostBridge																		
00011b-00001b	Reserved																		
00100b	Core. See CoreId.																		
00111b-00101b	Reserved																		
01000b	Memory controller.																		
01111b-01001b	Reserved																		
1HH0Nb	Link. Link HH; sublink N (where N=0b for ganged links)																		
1xx1xb	Reserved																		
14:11	SrqEntryState. Records unspecified, implementation-specific information (all zeroes means idle).																		
10:7	OpType. Records unspecified, implementation-specific information.																		
6:1	LinkCommand. When the NB WDT expires, the link command of the transaction that timed out is captured here. This field is encoded identically to the “Code” field for link transactions defined in the link specification.																		
0	Reserved																		

MSR0000_0413 NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)

[MSR0000_0413](#) is the first of the NB machine check miscellaneous registers. [MSR0000_0413](#) is associated with the DRAM error type. To see the remaining NB machine check miscellaneous registers, refer to [MSRC000_0408](#). Only one of these registers exists in multi-core devices; see 3.1.1 [Northbridge MSRs In Multi-Core Products].

Bits	Description										
63	Valid. IF (MSRC001_0015[McStatusWrEn]) THEN Read-write. ELSE Read-only. ENDIF. Reset: 1. 1=The CntP field is present.										
62	CntP: counter present. IF (MSRC001_0015[McStatusWrEn]) THEN Read-write. ELSE Read-only. ENDIF. Reset: 1. 1=A valid threshold counter is present.										
61	Locked. IF (MSRC001_0015[McStatusWrEn]) THEN Read-write. ELSE Read-only. ENDIF. Reset: 0. BIOS: IF (IntType==10b) THEN 1. ELSE 0. ENDIF. 1=Writes to bits [55:32] of this register are ignored. Set by BIOS to indicate that this register is not available for OS use. When MSRC001_0015[McStatusWrEn] is set, MSR writes to this register update all bits, regardless of the state of the Locked bit.										
60:56	Reserved.										
55:52	LvtOffset: LVT offset. IF (MSRC001_0015[McStatusWrEn] ~MSR0000_0413[Locked]) THEN Read-write. ELSE Read-only. ENDIF. Reset: 0h. Specifies the address of the LVT entry in the APIC registers as follows: LVT address = (LvtOffset shifted left 4 bits) + 500h (see APIC[530:500]). <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>3h-0h</td> <td>See APIC[530:500].</td> </tr> <tr> <td>Fh-4h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	3h-0h	See APIC[530:500].	Fh-4h	Reserved				
Bits	Description										
3h-0h	See APIC[530:500].										
Fh-4h	Reserved										
51	CntEn: counter enable. IF (MSRC001_0015[McStatusWrEn] ~Locked) THEN Read-write. ELSE Read-only. ENDIF. Reset: 0. 1=Count thresholding errors; see 2.15.1.7 [Error Thresholding].										
50:49	IntType: interrupt type. IF (MSRC001_0015[McStatusWrEn] ~MSR0000_0413[Locked]) THEN Read-write. ELSE Read-only. ENDIF. Cold reset: 0. Specifies the type of interrupt signaled when Ovrflw is set. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No Interrupt.</td> </tr> <tr> <td>01b</td> <td>APIC. APIC based interrupt (see LvtOffset above) to all cores.</td> </tr> <tr> <td>10b</td> <td>SMI. SMI trigger event (always routed to CpuCoreNum 0, as defined in 2.4.3 [Processor Cores and Downcoring]); see 2.4.8.2.3 [SMI Sources And Delivery].</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	No Interrupt.	01b	APIC. APIC based interrupt (see LvtOffset above) to all cores.	10b	SMI. SMI trigger event (always routed to CpuCoreNum 0, as defined in 2.4.3 [Processor Cores and Downcoring]); see 2.4.8.2.3 [SMI Sources And Delivery].	11b	Reserved
Bits	Description										
00b	No Interrupt.										
01b	APIC. APIC based interrupt (see LvtOffset above) to all cores.										
10b	SMI. SMI trigger event (always routed to CpuCoreNum 0, as defined in 2.4.3 [Processor Cores and Downcoring]); see 2.4.8.2.3 [SMI Sources And Delivery].										
11b	Reserved										
48	Ovrflw: overflow. IF (MSRC001_0015[McStatusWrEn] ~Locked) THEN Read-write; set-by-hardware. ELSE Read-only; set-by-hardware. ENDIF. Cold reset: 0. Set by hardware when ErrCnt transitions from FFEh to FFFh; also set by hardware if ErrCnt is initialized to FFFh and transitions from FFFh to 000h. When this bit is set, ErrCnt no longer increments and the interrupt selected by the IntType field is generated.										
47:44	Reserved.										
43:32	ErrCnt: error counter. IF (MSRC001_0015[McStatusWrEn] ~Locked) THEN Read-write; updated-by-hardware. ELSE Read-only; updated-by-hardware. ENDIF. Cold reset: 0. Written by software to set the starting value of the error counter. Incremented by hardware when errors are logged. Saturates at FFFh (no rollover). The threshold value, written by software, is (FFFh - the desired error count (the number of errors necessary in order for an interrupt to be taken)). The desired error count of 0 or a write value of FFFh is not supported.										
31:24	BlkPtr: Block pointer for additional MISC registers. Read-only. Value: 01h. 01h=Extended MC4_MISC MSR block is valid. See MSRC000_0408.										
23:0	Reserved.										

MSR0000_0414 EX Machine Check Control (MC5_CTL)

Read-write. Reset: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture]. See MSRC001_0049 [EX Machine Check Control Mask (MC5_CTL_MASK)].

Bits	Description
63:13	Unused.
12	DE: DE error.
11	FRF: flag register file parity.
10	AG1PRF: physical register file AG1 port parity.
9	AG0PRF: physical register file AG0 port parity.
8	EX1PRF: physical register file EX1 port parity.
7	EX0PRF: physical register file EX0 port parity.
6	MAP: mapper checkpoint array parity.
5	RETDISP: retire dispatch queue parity.
4	IDF: IDRF array parity.
3	PLDEX: EX payload array parity.
2	PLDAG: AG payload array parity.
1	PICWAK: wakeup array dest tag parity.
0	Unused.

MSR0000_0415 EX Machine Check Status (MC5_STATUS)

Cold reset: 0. See 2.15.1 [Machine Check Architecture]. See MSRC001_0015[McStatusWrEn]. Table 230 describes each error type. Table 231 describes the error codes and status register settings for each error type.

Bits	Description
63	Val: valid. See: MSR0000_0401[Val] .
62	Overflow: error overflow. See: MSR0000_0401[Overflow] .
61	UC: error uncorrected. See: MSR0000_0401[UC] .
60	En: error enable. See: MSR0000_0401[En] .
59	MiscV: miscellaneous error register valid. See: MSR0000_0401[MiscV] . 1=Valid thresholding in MSR0000_0417 .
58	AddrV: error address valid. See: MSR0000_0401[AddrV] .
57	PCC: processor context corrupt. See: MSR0000_0401[PCC] .
56:21	Reserved.
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. See MSR0000_0401[ErrorCodeExt] for description. See Table 231 for values.
15:0	ErrorCode: error code. See: MSR0000_0401[ErrorCode] .

Table 230: EX Error Descriptions

Error Type	Error Sub-type	Description ²	CTL ¹	EAC ³
WDT error	-	The WDT timer has expired.	WDT	E
Internal	Wakeup array dest tag parity	A parity error occurred in the wakeup array.	PICWAK	D
	AG payload array parity	A parity error occurred in the address generator payload array.	PLDAG	D
	EX payload array parity	A parity error occurred in the EX payload array.	PLDEX	D
	IDRF array parity	A parity error occurred in the immediate displacement register file.	IDF	D
	Retire dispatch queue parity	A parity error occurred in the retire dispatch queue. This error causes the processor to enter the Shutdown state; 2.15.1.3.1 [MCA conditions that cause Shutdown] .	RETDISP	E
	Mapper checkpoint array parity	A parity error occurred in the mapper checkpoint array. This error causes the processor to enter the Shutdown state if UC=1; 2.15.1.3.1 [MCA conditions that cause Shutdown] .	MAP	D
	EX0PRF parity	A parity error occurred in the physical register file's EX0 port.	EX0PRF	D
	EX1PRF parity	A parity error occurred in the physical register file's EX1 port.	EX1PRF	D
	AG0PRF parity	A parity error occurred in the physical register file's AG0 port.	AG0PRF	D
	AG1PRF parity	A parity error occurred in the physical register file's AG1 port.	AG1PRF	D
	Flag register file parity	A parity error occurred in the flag register file.	FRF	D
DE error	A DE error occurred.	DE	E	
<ol style="list-style-type: none"> See MSR0000_0414. CID: core ID. All EX errors are reported to the affected core; see 2.15.1.3. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See 2.15.1.3. 				

Table 231: EX Error Signatures

Error Type	Error Sub-Type	Error Code Ext	Error Code						UC	ADDRV	PCC
			Type	PP	T	RRRR	II	LL			
WDT error	-	00000b	BUS	GEN	1	GEN	GEN	LG	1	1	1
Internal error	Wakeup array dest tag parity	00001b			0				1/0	0	0
	AG payload array parity	00010b							1/0	1	0
	EX payload array parity	00011b							1/0	1	0
	IDRF array parity	00100b							1/0	1	0
	Retire dispatch queue parity	00101b							1 ¹	1	0
	Mapper checkpoint array parity	00110b							1/0 ²	1	0
	EX0PRF parity	00111b							1/0	0	0
	EX1PRF parity	01000b							1/0	0	0
	AG0PRF parity	01001b							1/0	0	0
	AG1PRF parity	01010b							1/0	0	0
	Flag register file parity	01011b							1/0	0	0
	DE error	01100b							1/0	0	0

1. Causes shutdown.

2. May cause shutdown if UC=1.

MSR0000_0416 EX Machine Check Address (MC5_ADDR)

Read-write; updated-by-hardware. Cold reset: 0. The MCI_ADDR register contains valid data if indicated by MCI_STATUS[AddrV]. See 2.15.1 [Machine Check Architecture]. The register format depends on the type of error being logged:

Bits	Description
63:0	ADDR. See Table 232.

Table 232: EX Address Register

Error Type	Error Sub-Type	Address Register Bits	Description
WDT	-	47:0	Logical address of the next instruction after the last instruction retired.

Table 232: EX Address Register

Error Type	Error Sub-Type	Address Register Bits	Description
Internal	DE correction error		
	AG payload array parity	5:0	Scheduler QID
	EX payload array parity	5:0	Scheduler QID
	IDRF array parity	5:0	Scheduler QID
	Retire dispatch queue parity	6:0	Retirement ID
	Mapper checkpoint array parity	5:0	Checkpoint ID

MSR0000_0417 EX Machine Check Miscellaneous (MC5_MISC)

See 2.15.1.7 [Error Thresholding].

Bits	Description
63	Valid. See: MSR0000_0403[Valid] .
62	CntP: counter present. See: MSR0000_0403[CntP] .
61:52	Reserved.
51	CntEn: counter enable. See: MSR0000_0403[CntEn] .
50:49	Reserved.
48	Ovrflw: overflow. See: MSR0000_0403[Ovrflw] .
47:44	Reserved.
43:32	ErrCnt: error counter. See: MSR0000_0403[ErrCnt] .
31:24	BlkPtr: Block pointer for additional MISC registers. See: MSR0000_0403[BlkPtr] .
23:0	Reserved.

MSR0000_0418 FP Machine Check Control (MC6_CTL)

SharedC; Read-write. Reset: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture]. See [MSRC001_004A \[FP Machine Check Control Mask \(MC6_CTL_MASK\)\]](#).

Bits	Description
63:6	Unused.
5	SRF: status register file parity error.
4	RetireQ: retire queue parity error.
3	Unused.
2	Sched: scheduler table parity error.
1	FreeList: free list parity error.
0	PRF: physical register file parity error.

MSR0000_0419 FP Machine Check Status (MC6_STATUS)

Cold reset: 0. See 2.15.1 [Machine Check Architecture]. See MSRC001_0015[McStatusWrEn]. Table 233 describes each error type. Table 234 describes the error codes and status register settings for each error type.

Bits	Description
63	Val: valid. See: MSR0000_0401[Val] .
62	Overflow: error overflow. See: MSR0000_0401[Overflow] .
61	UC: error uncorrected. See: MSR0000_0401[UC] .
60	En: error enable. See: MSR0000_0401[En] .
59	MiscV: miscellaneous error register valid. Read-only. See MSR0000_0401[MiscV] .
58	AddrV: error address valid. Read-only. See MSR0000_0401[AddrV] .
57	PCC: processor context corrupt. See: MSR0000_0401[PCC] .
56:21	Reserved.
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. See MSR0000_0401[ErrorCodeExt] for description. See Table 234 for values.
15:0	ErrorCode: error code. See: MSR0000_0401[ErrorCode] .

Table 233: FP Error Descriptions

Error Type	Error Sub-Type	Description ²	CTL ⁴	CID ³	EAC ¹
Floating Point Unit	Physical Register File	A parity error occurred in the Physical Register File (PRF).	PRF	0	E
	Status Register File	A parity error occurred in the Status Register File (SRF).	SRF	A	E
	Free List	A parity error occurred on the Free List.	FreeList	0	E
	Retire Queue	A parity error occurred in the Retire Queue.	RetireQ	0	E
	Scheduler	A parity error occurred in the Scheduler table.	Sched	0	E
<ol style="list-style-type: none"> 1. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See 2.15.1.3. 2. All FP errors are system fatal and result in a sync flood. 3. CID: core ID. A=Error reported to the affected core. 0=Error reported to core 0 of the compute unit; see 2.15.1.3. 4. See MSR0000_0418. 					

Table 234: FP Error Signatures

Error Type	Error Sub-Type	Error Code-Ext	Error Code						UC	ADDRV	PCC
			Type	PP	T	RRRR	II	LL			
Floating Point Unit	Status Register File	00101b	BUS	GEN	0	GEN	GEN	LG	1	0	1
	Physical Register File	00010b									
	Free List	00001b									
	Retire Queue	00011b									
	Scheduler	00100b									

MSR0000_041A FP Machine Check Address (MC6_ADDR)

Reset: 0000_0000_0000_0000h. Read-only. See 2.15.1 [Machine Check Architecture].

Bits	Description
63:0	Reserved.

MSR0000_041B FP Machine Check Miscellaneous (MC6_MISC)

Reset: 0000_0000_0000_0000h. Read-only. See 2.15.1 [Machine Check Architecture].

Bits	Description
63:0	Reserved.

3.19 MSRs - MSRC000_0xxx

MSRC000_0080 Extended Feature Enable (EFER)

Reset: 0000_0000_0000_0000h. SKINIT Execution: 0000_0000_0000_0000h.

Bits	Description
63:16	MBZ.
15	TCE: translation cache extension enable. Read-write. 1=Translation cache extension is enabled.
14	FFXSE: fast FXSAVE/FRSTOR enable. Read-write. 1=Enables the fast FXSAVE/FRSTOR mechanism. A 64-bit operating system uses CPUID Fn0000_0001_EDX[24] to determine the presence of this feature before enabling it. This bit is set once by the operating system and its value is not changed afterwards.
13	LMSLE: long mode segment limit enable. Read-write. 1=Enables the long mode segment limit check mechanism.
12	SVME: secure virtual machine (SVM) enable. Read-write. 1=SVM features are enabled.
11	NXE: no-execute page enable. Read-write. 1=The no-execute page protection feature is enabled.
10	LMA: long mode active. Read-only. 1=Indicates that long mode is active.
9	MBZ.
8	LME: long mode enable. Read-write. 1=Long mode is enabled.
7:1	RAZ.
0	SYSCALL: system call extension enable. Read-write. 1=SYSCALL and SYSRET instructions are enabled. This adds the SYSCALL and SYSRET instructions which can be used in flat addressed operating systems as low latency system calls and returns.

MSRC000_0081 SYSCALL Target Address (STAR)

Reset: 0000_0000_0000_0000h. This register holds the target address used by the SYSCALL instruction and the code and stack segment selector bases used by the SYSCALL and SYSRET instructions.

Bits	Description
63:48	SysRetSel: SYSRET CS and SS. Read-write.
47:32	SysCallSel: SYSCALL CS and SS. Read-write.
31:0	Target: SYSCALL target address. Read-write.

MSRC000_0082 Long Mode SYSCALL Target Address (STAR64)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	LSTAR: long mode target address. Read-write. Target address for 64-bit mode calling programs. The address stored in this register must be in canonical form (if not canonical, a #GP fault occurs).

MSRC000_0083 Compatibility Mode SYSCALL Target Address (STARCOMPAT)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	CSTAR: compatibility mode target address. Read-write. Target address for compatibility mode. The address stored in this register must be in canonical form (if not canonical, a #GP fault occurs).

MSRC000_0084 SYSCALL Flag Mask (SYSCALL_FLAG_MASK)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:32	RAZ.
31:0	Mask: SYSCALL flag mask. Read-write. This register holds the EFLAGS mask used by the SYSCALL instruction. 1=Clear the corresponding EFLAGS bit when executing the SYSCALL instruction.

MSRC000_00E7 Read-Only Max Performance Frequency Clock Count (MPerfReadOnly)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	MPerfReadOnly: read-only maximum core clocks counter. IF (MSRC001_0015[EffFreqReadOnlyLock]) THEN Read-only; Updated-by-hardware. ELSE Read-write; Updated-by-hardware. ENDIF. Incremented by hardware at the P0 frequency while the core is in C0. This register does not increment when the core is in the stop-grant state. In combination with MSRC000_00E8, this is used to determine the effective frequency of the core. This field uses software P-state numbering. See MSRC001_0015[EffFreqCntMwait], 2.5.3.3 [Effective Frequency], and 2.5.3.1.2.1 [Software P-state Numbering]. This register is not affected by writes to MSR0000_00E7.

MSRC000_00E8 Read-Only Actual Performance Frequency Clock Count (APerfReadOnly)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	APerfReadOnly: read-only actual core clocks counter. IF (MSRC001_0015[EffFreqReadOnlyLock]) THEN Read-only; Updated-by-hardware. ELSE Read-write; Updated-by-hardware. ENDIF. This register increments in proportion to the actual number of core clocks cycles while the core is in C0. The register does not increment when the core is in the stop-grant state. See MSRC000_00E7. This register is not affected by writes to MSR0000_00E8.

MSRC000_0100 FS Base (FS_BASE)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	FSBase: expanded FS segment base. Read-write. This register provides access to the expanded 64-bit FS segment base. The address stored in this register must be in canonical form (if not canonical, a #GP fault fill occurs).

MSRC000_0101 GS Base (GS_BASE)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	GSBase: expanded GS segment base. Read-write; Updated-by-hardware . This register provides access to the expanded 64-bit GS segment base. The address stored in this register must be in canonical form (if not canonical, a #GP fault fill occurs).

MSRC000_0102 Kernel GS Base (KernelGSbase)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	KernelGSBase: kernel data structure pointer. Read-write. This register holds the kernel data structure pointer which can be swapped with the GS_BASE register using the SwapGS instruction. The address stored in this register must be in canonical form (if not canonical, a #GP fault occurs).

MSRC000_0103 Auxiliary Time Stamp Counter (TSC_AUX)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:32	Reserved.
31:0	TscAux: auxiliary time stamp counter data. Read-write. It is expected that this is initialized by privileged software to a meaningful value, such as a processor ID. This value is returned in the RDTSCP instruction.

MSRC000_0104 Time Stamp Counter Ratio (TscRateMsr)

Reset: 0000_0001_0000_0000h. This register defines the ratio of the desired TSC frequency to the core's maximum (i.e. P(0)) frequency. If the P(0) frequency is changed, then the TSC Frequency still uses the TSC Ratio MSR value now applied to the new P(0) frequency. The $TSCFreq = P(0) \text{ frequency} * TSCRatio$, so $TSCRatio = (\text{Desired } TSCFreq) / P(0) \text{ frequency}$.

Bits	Description
63:40	MBZ.
39:32	TscRateMsrInt: time stamp counter rate integer. Read-write. Specifies the integer part of the MSR TSC ratio value.
31:0	TscRateMsrFrac: time stamp counter rate fraction. Read-write. Specifies the fractional part of the MSR TSC ratio value.

MSRC000_0105 Lightweight Profile Configuration (LWP_CFG)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:48	MBZ.
47:40	LwpVector: threshold interrupt vector. Read-write. Interrupt vector number used by LWP Threshold interrupts. Must be provided if LwpInt is set to 1.
39:32	LwpCoreId: core ID. Read-write. Core identification stored into the trace record. BIOS: CPUID Fn0000_0001_EBX[LocalApicId] . Software is recommended to set this to CPUID Fn0000_0001_EBX[LocalApicId] .
31	LwpInt: interrupt on threshold overflow. Read-write. 1=Enable LWP to interrupt on threshold overflow. CPUID Fn8000_001C_EAX[LwpInt] is an alias of MSRC000_0105[LwpInt] .
30:7	MBZ.
6	LwpRNH: core reference clocks not halted event support. MBZ. 1=Enable LWP to count core reference clocks not halted. CPUID Fn8000_001C_EAX[LwpRNH] is an alias of MSRC000_0105[LwpRNH] .
5	LwpCNH: core clocks not halted event support. MBZ. 1=Enable LWP to count core clocks not halted. CPUID Fn8000_001C_EAX[LwpCNH] is an alias of MSRC000_0105[LwpCNH] .
4	LwpDME: DC miss event support. MBZ. 1=Enable LWP to count DC misses. CPUID Fn8000_001C_EAX[LwpDME] is an alias of MSRC000_0105[LwpDME] .
3	LwpBRE: branch retired event support. Read-write. 1=Enable LWP to count branches retired. CPUID Fn8000_001C_EAX[LwpBRE] is an alias of MSRC000_0105[LwpBRE] .
2	LwpIRE: instructions retired event support. Read-write. 1=Enable LWP to count instructions retired. CPUID Fn8000_001C_EAX[LwpIRE] is an alias of MSRC000_0105[LwpIRE] .
1	LwpVAL: LWPVAL instruction support. Read-write. 1=LWPVAL instruction is enabled. CPUID Fn8000_001C_EAX[LwpVAL] is an alias of MSRC000_0105[LwpVAL] .
0	Reserved.

MSRC000_0106 Lightweight Profile Control Block Address (LWP_CBADDR)

Access to the internal copy of the LWPCB logical line/64 B address. A read returns the current LWPCB address without performing any of the operations described for the SLWPCB instruction. A write to this register with a non-zero value causes a #GP Fault. Use LLWPCB or XRSTOR to load an LWPCB address. Writing a zero to LWP_CBADDR will immediately disable LWP, discarding any internal state. For instance, an operating system can write a zero to stop LWP when it terminates a thread. All references to the LWPCB implicitly use the DS segment register. Must be 64 B aligned.

Bits	Description
63:6	LwpCbAddr[63:6]: control block logical address. Read-write. Reset: 0. LwpCbAddr[63:0] = {LwpCbAddr[63:6], 000000b}.
5:0	MBZ.

MSRC000_0408 Machine Check Misc 4 (Link Thresholding) 1 (MC4_MISC1)

Not-same-for-all. [MSRC000_0408](#) is associated with the link error type. [MSRC000_0408](#) are the block of extended NB machine check miscellaneous registers. Only one of these registers exists in multi-core devices; see 3.1.1 [Northbridge MSRs In Multi-Core Products]. Accessibility of this register by non-NBC cores is affected by [D18F3x44\[NbMcaToMstCpuEn\]](#).

Bits	Description
63:32	See: MSR0000_0413 [63:32].
31:0	Reserved.

MSRC000_040[F:9] Reserved

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	RAZ.

3.20 MSRs - MSRC001_0xxx

MSRC001_00[03:00] Performance Event Select (PERF_CTL[3:0])

The legacy alias of MSRC001_020[6,4,2,0]. See [MSRC001_020\[A,8,6,4,2,0\]](#).

Table 235: Register Mapping for MSRC001_00[03:00]

Register	Function
MSRC001_0000	Counter 0
MSRC001_0001	Counter 1
MSRC001_0002	Counter 2
MSRC001_0003	Counter 3

Bits	Description
63:0	MSRC001_00[03:00] is an alias of MSRC001_020[6,4,2,0].

MSRC001_00[07:04] Performance Event Counter (PERF_CTR[3:0])

The legacy alias of MSRC001_020[7,5,3,1]. See [MSRC001_020\[B,9,7,5,3,1\]](#).

Table 236: Register Mapping for MSRC001_00[07:04]

Register	Function
MSRC001_0004	Counter 0
MSRC001_0005	Counter 1
MSRC001_0006	Counter 2
MSRC001_0007	Counter 3

Bits	Description
63:0	MSRC001_00[07:04] is an alias of MSRC001_020[7,5,3,1].

MSRC001_0010 System Configuration (SYS_CFG)

SharedC.

Bits	Description
63:23	Reserved.
22	Tom2ForceMemTypeWB: top of memory 2 memory type write back. Read-write; SharedC. Reset: 0. 1=The default memory type of memory between 4GB and TOM2 is write back instead of the memory type defined by MSR0000_02FF [MTRR Default Memory Type (MTRRdefType)] [MemType]. For this bit to have any effect, MSR0000_02FF[MtrrDefTypeEn] must be 1. MTRRs and PAT can be used to override this memory type.
21	MtrrTom2En: MTRR top of memory 2 enable. Read-write; SharedC. Reset: 0. 0= MSRC001_001D [Top Of Memory 2 (TOM2)] is disabled. 1=This register is enabled. See D0F0x64_x19[TomEn] .

20	MtrrVarDramEn: MTRR variable DRAM enable. Read-write; SharedC . Reset: 0. BIOS: 1. 0= MSRC001_001A [Top Of Memory (TOP_MEM)] and IORRs are disabled. 1=These registers are enabled.
19	MtrrFixDramModEn: MTRR fixed RdDram and WrDram modification enable. Read-write. Reset: 0. Controls access to MSR0000_02[6F:68,59:58,50][RdDram, WrDram] . BIOS: This bit should be set to 1 during BIOS initialization of the fixed MTRRs, then cleared to 0 for operation. 0=Access type is MBZ; writing 00b does not change the hidden value of MSR0000_02[6F:68,59:58,50][RdDram, WrDram] . 1=Access type is Read-write.
18	MtrrFixDramEn: MTRR fixed RdDram and WrDram attributes enable. Read-write; SharedC . Reset: 0. BIOS: 1. 1=Enables the RdDram and WrDram attributes in MSR0000_02[6F:68,59:58,50] .
17	Reserved.
16	ChgToDirtyDis: change to dirty disable. Read-write; SharedC . Reset: 0. 1=Disables Change-to-Dirty command; The change-to-dirty condition is handled by evicting the line and then fetching it with a RdBkM command.
15:0	Reserved.

MSRC001_0015 Hardware Configuration (HWCR)

Bits	Description
63:28	Reserved.
27	EffFreqReadOnlyLock: read-only effective frequency counter lock. Write-1-only. Reset: 0. BIOS: 1. 1= MSRC000_00E7 and MSRC000_00E8 are read-only.
26	EffFreqCntMwait: effective frequency counting during mwait. Read-write. Reset: 0. Specifies whether MSR0000_00E7 [Max Performance Frequency Clock Count (MPERF)] and MSR0000_00E8 [Actual Performance Frequency Clock Count (APERF)] increment while the core is in the monitor event pending state. 0=The registers do not increment. 1=The registers increment. See 2.5.3.3 [Effective Frequency] .
25	CpbDis: core performance boost disable. Read-write. Reset: 0. Specifies whether core performance boost is requested to be enabled or disabled. 0=CPB is requested to be enabled. 1=CPB is disabled. See 2.5.3.1.1 [Application Power Management (APM)] . If core performance boost is disabled while a core is in a boosted P-state, the core will automatically transition to the highest performance non-boosted P-state.
24	TscFreqSel: TSC frequency select. Read-only. Reset: 1. 1=The TSC increments at the P0 frequency. This field uses software P-state numbering. See 2.5.3.1.2.1 [Software P-state Numbering] .
23	ForceRdWrSzPrb: force probes for upstream RdSized and WrSized. Read-write. Reset: 0. This bit is shared between all processor cores. 1=Forces probes on read-sized and write-sized transactions: <ul style="list-style-type: none"> • From the core. • From IO, upstream: <ul style="list-style-type: none"> • Except operations that are “Isoc=1, PassPW=1, RespPassPW=1, Coherent=1”. See 2.8.2.1.1 [DRAM and MMIO Memory Space] .
22:21	Reserved.

20	IoCfgGpFault: IO-space configuration causes a GP fault. Read-write. Reset: 0. 1=IO-space accesses to configuration space cause a GP fault. The fault is triggered if any part of the IO read/write address range is between CF8h and CFFh, inclusive. These faults only result from single IO instructions, not to string and REP IO instructions. This fault takes priority over the IO trap mechanism described by MSRC001_0054 [IO Trap Control (SMI_ON_IO_TRAP_CTL_STS)] .
19	Reserved.
18	McStatusWrEn: machine check status write enable. Read-write. Reset: 0. McStatusWrEn can be used to debug machine check exception and interrupt handlers. See 2.15.2 [Error Injection and Simulation] . See 2.15.1 [Machine Check Architecture] . <ul style="list-style-type: none"> • 1=MCi_STATUS registers are read-write, including reserved fields; do not cause general protection faults; such writes update all implemented bits in these registers; All fields of all threshold registers are Read-write when accessed from MSR space, including Locked, except BlkPtr which is always read-only; McStatusWrEn does not change the access type for the thresholding registers accessed via configuration space. • 0=MCi_STATUS registers are readable; writing a non-zero pattern to these registers causes a general protection fault. • The MCi_STATUS registers are: MSR0000_0401, MSR0000_0405, MSR0000_0409, MSR0000_040D, MSR0000_0411, MSR0000_0415, MSR0000_0419. McStatusWrEn does not affect the writability of MSR0000_0001; MSR0000_0001 is always writable. • The thresholding registers affected by McStatusWrEn are: MSR0000_0403, MSR0000_0407, MSR0000_040B, MSR0000_0413, MSR0000_0417, MSRC000_0408.
17	Wrap32Dis: 32-bit address wrap disable. Read-write. Reset: 0. 1=Disable 32-bit address wrapping. Software can use Wrap32Dis to access physical memory above 4 Gbytes without switching into 64-bit mode. To do so, software should write a greater-than 4 Gbyte address to MSRC000_0100 [FS Base (FS_BASE)] and MSRC000_0101 [GS Base (GS_BASE)] . Then it would address ± 2 Gbytes from one of those bases using normal memory reference instructions with a FS or GS override prefix. However, the INVLPG, FST, and SSE store instructions generate 32-bit addresses in legacy mode, regardless of the state of Wrap32Dis.
16:15	Reserved.
14	RsmSpCycDis: RSM special bus cycle disable. IF MSRC001_0015[SmmLock] THEN Read-only ELSE Read-write ENDIF. Reset: 0. 0=A link special bus cycle, SMIACK, is generated on a resume from SMI.
13	SmiSpCycDis: SMI special bus cycle disable. IF MSRC001_0015[SmmLock] THEN Read-only ELSE Read-write ENDIF. Reset: 0. 0=A link special bus cycle, SMIACK, is generated when an SMI interrupt is taken.
12	HltXSpCycEn: halt-exit special bus cycle enable. Read-write. Reset: 0. BIOS: 1. Read-write. Specifies whether any messages are sent to the FCH when a core enters or exits a C-state. 1=Messages are sent. 0=Messages are not sent. See 2.5.3.2.4.1 [FCH Messaging] .
11	Reserved.
10	MonMwaitUserEn: MONITOR/MWAIT user mode enable. Read-write. Reset: 0. 1=The MONITOR and MWAIT instructions are supported in all privilege levels. 0=The MONITOR and MWAIT instructions are supported only in privilege level 0; these instructions in privilege levels 1 to 3 cause a #UD exception. The state of this bit is ignored if MonMwaitDis is set.
9	MonMwaitDis: MONITOR and MWAIT disable. Read-write. Reset: 0. 1=The MONITOR and MWAIT opcodes become invalid. This affects what is reported back through CPUID Fn0000_0001_ECX[Monitor] .

8	IgnneEm: IGNNE port emulation enable. Read-write. Reset: 0. 1=Enable emulation of IGNNE port.
7:6	Reserved.
5	Reserved.
4	INVDWBINVD: INVD to WBINVD conversion. Read-write. Reset: 1. 1=Convert INVD to WBINVD. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot]. This bit is required to be set when both cores of a compute unit are enabled, and thus share the L2 cache.
3	TlbCacheDis: cacheable memory disable. Read-write. Reset: 0. 1=Disable performance improvement that assumes that the PML4, PDP, PDE and PTE entries are in cacheable WB DRAM. Operating systems that maintain page tables in any other memory type must set the TlbCacheDis bit to insure proper operation. <ul style="list-style-type: none"> TlbCacheDis does not override the memory type specified by the SMM ASeg and TSeg memory regions. See 2.4.8.2.7 [The Protected ASeg and TSeg Areas].
2	Reserved.
1	Reserved.
0	SmmLock: SMM code lock. Read; write-1-only. Reset: 0. SBIOS: 1. 1=SMM code in the ASeg and TSeg range and the SMM registers are read-only and SMI interrupts are not intercepted in SVM.

MSRC001_00[18,16] IO Range Base (IORR_BASE[1:0])

SharedC. Reset: X. MSRC001_0016 and MSRC001_0017 combine to specify the first IORR range and MSRC001_0018 and MSRC001_0019 combine to specify the second IORR range. A core access, with address CPUAddr, is determined to be within IORR address range if the following equation is true:

$CPUAddr[47:12] \& PhyMask[47:12] == PhyBase[47:12] \& PhyMask[47:12]$.

BIOS can use the IORRs to create an IO hole within a range of addresses that would normally be mapped to DRAM. It can also use the IORRs to re-assert a DRAM destination for a range of addresses that fall within a bigger IO hole that overlays DRAM.

See 2.4.5.1.2 [Determining The Access Destination for Core Accesses].

Bits	Description
63:48	RAZ.
47:12	PhyBase: physical base address. Read-write.
11:5	RAZ.
4	RdMem: read from memory. Read-write. 1=Read accesses to the range are directed to system memory. 0=Read accesses to the range are directed to IO.
3	WrMem: write to memory. Read-write. 1=Write accesses to the range are directed to system memory. 0=Write accesses to the range are directed to IO.
2:0	RAZ.

MSRC001_00[19,17] IO Range Mask (IORR_MASK[1:0])

SharedC. Reset: X. See MSRC001_00[18,16].

Bits	Description
63:48	RAZ.
47:12	PhyMask: physical address mask. Read-write.
11	Valid. Read-write. 1=The pair of registers that specifies an IORR range is valid.
10:0	RAZ.

MSRC001_001A Top Of Memory (TOP_MEM)

SharedC. Reset: X.

Bits	Description
63:48	RAZ.
47:23	TOM[47:23]: top of memory. Read-write. Specifies the address that divides between MMIO and DRAM. This value is normally placed below 4G. From TOM to 4G is MMIO; below TOM is DRAM. See 2.4.5 [System Address Map].
22:0	RAZ.

MSRC001_001D Top Of Memory 2 (TOM2)

SharedC. Reset: X.

Bits	Description
63:48	RAZ.
47:23	TOM2[47:23]: second top of memory. Read-write. Specifies the address divides between MMIO and DRAM. This value is normally placed above 4G. From 4G to TOM2 - 1 is DRAM; TOM2 and above is MMIO. See 2.4.5 [System Address Map]. This register is enabled by MSRC001_0010 [System Configuration (SYS_CFG)][MtrrTom2En].
22:0	RAZ.

MSRC001_001F Northbridge Configuration 1 (NB_CFG1)

Read-write. Only one of these registers exists in multi-core devices; see 3.1.1 [Northbridge MSRs In Multi-Core Products]. MSRC001_001F[31:0] is an alias of D18F3x88. MSRC001_001F[63:32] is an alias of D18F3x8C.

Bits	Description
63:55	Reserved.
54	InitApicIdCpuIdLo. Reset: 0. BIOS: 1. 0=Reserved. 1=Selects the format for ApicId; see APIC20. InitApicIdCpuIdLo must be set before D18F0x60[NodeId] is programmed.
53:51	Reserved.
50	DisOrderRdRsp. Reset: 0. 1=Disables ordered responses to IO link read requests.
49:47	Reserved.
46	EnableCf8ExtCfg: enable CF8 extended configuration cycles. Reset: 0. 1=Allows the IO configuration space access method, IOCF8 and IOCF8C, to be used to generate extended configuration cycles by enabling IOCF8[27:24].

45	DisUsSysMgtReqToNcHt: disable upstream system management request to link. Reset: 0. 1=Disables downstream reflection of upstream STPCLK and x86 legacy input system management commands (in order to work around potential deadlock scenarios related to reflection regions).
44:37	Reserved.
36	DisDatMsk: disable data mask. Reset: 0. 1=Disables DRAM data masking function; all write requests that are less than one cacheline, a DRAM read is performed before writing the data.
35:32	Reserved.
31	DisCohLdtCfg: disable coherent link configuration accesses. Reset: 0. 1=Disables automatic routing of PCI configuration accesses to the processor configuration registers; PCI configuration space accesses which fall within the hard-coded range reserved for processor configuration-space registers are instead routed to the IO link specified by D18F1x[EC:E0] [Configuration Map] . This can be used to effectively hide the configuration registers from software. It can also be used to provide a means for an external chip to route processor configuration accesses according to a scheme other than the hard-coded version. When used, this bit needs to be set on all processors in a system. PCI configuration accesses should not be generated if this bit is not set on all processors.
30:0	Reserved.

MSRC001_0022 Machine Check Exception Redirection

Reset: 0000_0000_0000_0000h. This register can be used to redirect machine check exceptions (MCEs) to SMIs or vectored interrupts. If both RedirSmiEn and RedirVecEn are set, then undefined behavior results.

Bits	Description
63:10	Reserved.
9	RedirSmiEn. Read-write. 1=Redirect MCEs (that are directed to this core) to generate an SMI-trigger IO cycle via MSRC001_0056 . The status is stored in SMMFEC4[MceRedirSts] .
8	RedirVecEn. Read-write. 1=Redirect MCEs (that are directed to this core) to generate a vectored interrupt, using the interrupt vector specified in RedirVector.
7:0	RedirVector. Read-write. See RedirVecEn.

MSRC001_00[35:30] Processor Name String

SharedNC. Reset: 0000_0000_0000_0000h. BIOS: [Table 238](#). These registers holds the CPUID name string in ASCII. The state of these registers are returned by CPUID instructions, [CPUID Fn8000_000\[4:2\]_E\[D,C,B,A\]X](#). BIOS should set these registers to the product name for the processor as provided by AMD. Each register contains a block of 8 ASCII characters; the least byte corresponds to the first ASCII character of the block; the most-significant byte corresponds to the last character of the block. MSRC001_0030 contains the first block of the name string; MSRC001_0035 contains the last block of the name string.

Table 237: Register Mapping for MSRC001_00[35:30]

Register	Function
MSRC001_0030	Characters 7-0
MSRC001_0031	Characters 15-8
MSRC001_0032	Characters 23-16

Table 237: Register Mapping for MSRC001_00[35:30]

MSRC001_0033	Characters 31-24
MSRC001_0034	Characters 39-32
MSRC001_0035	Characters 47-40

See [D18F5x194](#) for the access method to [D18F5x198_x\[B:0\]](#).

Table 238: BIOS recommendation for MSRC001_00[35:30]

Register	BIOS
MSRC001_0030	{D18F5x198_x1, D18F5x198_x0}
MSRC001_0031	{D18F5x198_x3, D18F5x198_x2}
MSRC001_0032	{D18F5x198_x5, D18F5x198_x4}
MSRC001_0033	{D18F5x198_x7, D18F5x198_x6}
MSRC001_0034	{D18F5x198_x9, D18F5x198_x8}
MSRC001_0035	{D18F5x198_xB, D18F5x198_xA}

Bits	Description
63:0	CpuNameString . Read-write.

MSRC001_003E Hardware Thermal Control (HTC)

Only one of these registers exists in multi-core devices; see [3.1.1 \[Northbridge MSRs In Multi-Core Products\]](#).

Bits	Description
63:32	Reserved.
31:0	Alias of D18F3x64 . Uses hardware P-state numbering. See 2.5.3.1.2.2 [Hardware P-state Numbering] .

MSRC001_0044 LS Machine Check Control Mask (MC0_CTL_MASK)

Read-write. Reset: 0000_0000_0000_0000h. BIOS: 0000_0000_0000_0000h. See [2.15.1 \[Machine Check Architecture\]](#). See [MSR0000_0400 \[LS Machine Check Control \(MC0_CTL\)\]](#).

Bits	Description
63:10	Reserved.
9	IntErrTyp1: internal error type 1.
8	IntErrTyp2: internal error type 2.
7	SRDE: read data errors. System read data errors on cache fill.
6	LFE: line fill error. Uncorrectable error on cache fill.
5	SCBP: SCB parity.
4	SQP: store queue parity.
3	LQP: load queue parity.
2	DatP: data parity.

1	TLBP: TLB parity.
0	TagP: tag parity.

MSRC001_0045 IF Machine Check Control Mask (MC1_CTL_MASK)

SharedC; read-write. Reset: 0000_0000_0000_0080h. BIOS: 0000_0000_0000_0080h. See 2.15.1 [Machine Check Architecture]. See MSR0000_0404 [IF Machine Check Control (MC1_CTL)].

Bits	Description
63:24	Reserved.
23	IVP: IC valid bit parity error.
22	L1TLBM: IC L1 TLB multi-match error.
21	L2TLBM: IC L2 TLB multi-match error.
20	DFIFOE: decoder FIFO parity error.
19	DPDBE: decoder predecode buffer parity error.
18	DEIBP: decoder instruction buffer parity error.
17	DEUOPQP: Decoder micro-op queue parity error.
16	DEPRP: decoder patch RAM parity error.
15	BSRP: branch status register parity error.
14	Reserved.
13	PQP: prediction queue parity error.
12	PFBP: prefetch buffer parity.
11:10	Reserved.
9	SRDE: system read data error.
8	Reserved.
7	LineFillPoison: line fill poison error.
6	L1TP: L1 TLB parity error.
5	L2TP: L2 TLB parity error.
4	ISTP: L1 cache probe tag array parity error.
3	IMTP: L1 cache main tag array parity error.
2	IDP: L1 cache data array parity error.
1:0	Reserved.

MSRC001_0046 CU Machine Check Control Mask (MC2_CTL_MASK)

SharedC; read-write. Reset: 0000_0000_0000_0000h. BIOS: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture]. See MSR0000_0408 [CU Machine Check Control (MC2_CTL)].

Bits	Description
63:15	Reserved.
14	L2TlbFill: TLB fill error enable. Data with uncorrectable error provided to TLB.
13	RdData: read data error from NB.

12	L2Tag: L2 cache tag error.
11	L2TlbData: L2 TLB parity error. Parity error reading from TLB.
10	L2Prefetch: L2 data prefetcher parity error.
9	XabAddr: XAB address parity error.
8	PrbAddr: probe buffer address parity error.
7	FillData: fill data parity and ECC error.
6	PrqAddr: post retire queue address parity error.
5	PrqData: post retire queue data parity error.
4	WccAddr: write coalescing cache address ECC error.
3	WccData: write coalescing cache data ECC error.
2	WebData: write combining buffer data parity error.
1	VbData: victim buffer data parity and ECC error.
0	L2TagMultiHit: L2 tag multiple hit error.

MSRC001_0047 Reserved (MC3_CTL_MASK)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	Reserved.

MSRC001_0048 NB Machine Check Control Mask (MC4_CTL_MASK)

Read-write; **Not-same-for-all**. Reset: 0000_0000_0400_0000h. BIOS: 0000_0000_0400_0000h. The format of MC4_CTL_MASK corresponds to MSR0000_0410 [NB Machine Check Control (MC4_CTL)]. For each defined bit position, 1=Disable logging. Only one of these registers exists in multi-core devices; see 3.1.1 [Northbridge MSRs In Multi-Core Products]. See 2.15.1 [Machine Check Architecture]. Accessibility of this register by non-NBC cores is affected by D18F3x44[NbMcaToMstCpuEn]. See MSR0000_0410 [NB Machine Check Control (MC4_CTL)].

Bits	Description
63:32	Reserved.
31	McaCpuDatErrEn.
30	Reserved.
29:27	Reserved.
26	NbArrayParEn.
25	UsPwDatErrEn.
24:18	Reserved.
17	CpPktDatEn.
16	NbIntProtEn.
15:13	Reserved.
12	WDTRptEn.
11	AtomicRMWEn.

10	Reserved.
9	TgtAbortEn.
8	MstrAbortEn.
7:6	Reserved.
5	SyncPktEn.
4:0	Reserved.

MSRC001_0049 EX Machine Check Control Mask (MC5_CTL_MASK)

Reset: 0000_0000_0000_0000h. BIOS: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture]. See [MSR0000_0414 \[EX Machine Check Control \(MC5_CTL\)\]](#).

Bits	Description
63:0	See: MSR0000_0414 . The format of MC5_CTL_MASK corresponds to MC5_CTL.

MSRC001_004A FP Machine Check Control Mask (MC6_CTL_MASK)

SharedC. Reset: 0000_0000_0000_0000h. BIOS: 0000_0000_0000_0000h. See 2.15.1 [Machine Check Architecture]. See [MSR0000_0418 \[FP Machine Check Control \(MC6_CTL\)\]](#).

Bits	Description
63:0	See: MSR0000_0418 . The format of MC6_CTL_MASK corresponds to MC6_CTL.

MSRC001_00[53:50] IO Trap (SMI_ON_IO_TRAP_[3:0])

SharedNC. Reset: 0000_0000_0000_0000h.

[MSRC001_00\[53:50\]](#) and [MSRC001_0054](#) provide a mechanism for executing the SMI handler if a an access to one of the specified addresses is detected. Access address and access type checking is performed before IO instruction execution. If the access address and access type match one of the specified IO address and access types, then: (1) the IO instruction is not executed; (2) any breakpoint, other than the single-step breakpoint, set on the IO instruction is not taken (the single-step breakpoint is taken after resuming from SMM); and (3) the SMI-trigger IO cycle specified by [MSRC001_0056](#). The status is stored in [SMMFEC4\[IoTrapSts\]](#).

IO-space configuration accesses are special IO accesses. An IO access is defined as an IO-space configuration access when IO instruction address bits[31:0] are CFCh, CFDh, CFEh, or CFFh when IO-space configuration is enabled ([IOCF8\[ConfigEn\]](#)). The access address for a configuration space access is the current value of [IOCF8\[BusNo, Device, Function, RegNo\]](#). The access address for an IO access that is not a configuration access is equivalent to the IO instruction address, bits[31:0].

The access address is compared with SmiAddr, and the instruction access type is compared with the enabled access types defined by ConfigSMI, SmiOnRdEn, and SmiOnWrEn. Access address bits[23:0] can be masked with SmiMask.

IO and configuration space trapping to SMI applies only to single IO instructions; it does not apply to string and REP IO instructions.

Table 239: Register Mapping for MSRC001_00[53:50]

Register	Function
MSRC001_0050	Range 0
MSRC001_0051	Range 1
MSRC001_0052	Range 2
MSRC001_0053	Range 3

Bits	Description
63	SmiOnRdEn: enable SMI on IO read. Read-write. 1=Enables SMI generation on a read access.
62	SmiOnWrEn: enable SMI on IO write. Read-write. 1=Enables SMI generation on a write access.
61	ConfigSmi: configuration space SMI. Read-write. 1=Configuration access. 0=IO access (that is not an IO-space configuration access).
60:56	Reserved.
55:32	SmiMask[23:0]. Read-write. SMI IO trap mask. 0=Mask address bit. 1=Do not mask address bit.
31:0	SmiAddr[31:0]. Read-write. SMI IO trap address.

MSRC001_0054 IO Trap Control (SMI_ON_IO_TRAP_CTL_STS)

SharedNC. Reset: 0000_0000_0000_0000h. For each of the SmiEn bits below, 1=The trap specified by the corresponding MSR is enabled. See [MSRC001_00\[53:50\]](#).

Bits	Description
63:32	RAZ.
31:16	Reserved.
15	IoTrapEn: IO trap enable. Read-write. 1=Enable IO and configuration space trapping specified by MSRC001_00[53:50] and MSRC001_0054 .
14:8	Reserved.
7	SmiEn3: SMI enable for the trap specified by MSRC001_0053. Read-write.
6	Reserved.
5	SmiEn2: SMI enable for the trap specified by MSRC001_0052. Read-write.
4	Reserved.
3	SmiEn1: SMI enable for the trap specified by MSRC001_0051. Read-write.
2	Reserved.
1	SmiEn0: SMI enable for the trap specified by MSRC001_0050. Read-write.
0	Reserved.

MSRC001_0055 Interrupt Pending

SharedNC. Reset: 0000_0000_0000_0000h.

Bits	Description
63:32	RAZ.
31	Reserved.
30	EnablePmTmrCheckLoop. Read-write. 1=The core loops on IO-space read accesses to the address specified by IOMsgAddr until the data value has incremented from the previous read access.
29:16	Reserved.
15:0	IOMsgAddr: IO message address. Read-write. IO space message address.

MSRC001_0056 SMI Trigger IO Cycle

Reset: 0000_0000_0000_0000h. See 2.4.8.2.3 [SMI Sources And Delivery]. This register specifies an IO cycle that may be generated when a local SMI trigger event occurs. If IoCycleEn is set and there is a local SMI trigger event, then the IO cycle generated is a byte read or write, based on IoRd, to address IoPortAddress. If the cycle is a write, then IoData contains the data written. If the cycle is a read, the value read is discarded. If IoCycleEn is clear and a local SMI trigger event occurs, then undefined behavior results.

Bits	Description
63:27	Reserved.
26	IoRd: IO Read. Read-write. 1=IO read; 0=IO write.
25	IoCycleEn: IO cycle enable. Read-write. 1=The SMI trigger IO cycle is enabled to be generated.
24	Reserved.
23:16	IoData. Read-write.
15:0	IoPortAddress. Read-write.

MSRC001_0058 MMIO Configuration Base Address

See 2.7 [Configuration Space] for a description of MMIO configuration space. All cores of all processors should be programmed with the same value of this register.

Bits	Description																								
63:48	RAZ.																								
47:20	MmioCfgBaseAddr[47:20]: MMIO configuration base address bits[47:20]. Read-write. Reset: X. Specifies the base address of the MMIO configuration range. The size of the MMIO configuration-space address range is specified by BusRange.																								
19:6	RAZ.																								
5:2	BusRange: bus range identifier. Read-write. Reset: X. Specifies the number of buses in the MMIO configuration space range. The size of the MMIO configuration space is 1 MB times the number of buses.																								
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>1</td> <td>5h</td> <td>32</td> </tr> <tr> <td>1h</td> <td>2</td> <td>6h</td> <td>64</td> </tr> <tr> <td>2h</td> <td>4</td> <td>7h</td> <td>128</td> </tr> <tr> <td>3h</td> <td>8</td> <td>8h</td> <td>256</td> </tr> <tr> <td>4h</td> <td>16</td> <td>Fh-9h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	0h	1	5h	32	1h	2	6h	64	2h	4	7h	128	3h	8	8h	256	4h	16	Fh-9h	Reserved
Bits	Description	Bits	Description																						
0h	1	5h	32																						
1h	2	6h	64																						
2h	4	7h	128																						
3h	8	8h	256																						
4h	16	Fh-9h	Reserved																						

1	Reserved.
0	Enable. Read-write. Reset: 0. 1=MMIO configuration space is enabled.

MSRC001_0061 P-state Current Limit

Read; GP-write; [SharedC](#); updated-by-hardware. See [2.5.3 \[CPU Power Management\]](#).

Bits	Description
63:7	RAZ.
6:4	PstateMaxVal: P-state maximum value. Specifies the lowest-performance non-boosted P-state (highest non-boosted value) allowed. Attempts to change MSRC001_0062[PstateCmd] to a lower-performance P-state (higher value) are clipped to the value of this field. This field uses software P-state numbering. See 2.5.3.1.2.1 [Software P-state Numbering] .
3	RAZ.
2:0	CurPstateLimit: current P-state limit. Specifies the highest-performance non-boosted P-state (lowest value) allowed. CurPstateLimit is always bounded by MSRC001_0061[PstateMaxVal] . Attempts to change the CurPstateLimit to a value greater (lower performance) than MSRC001_0061[PstateMaxVal] leaves CurPstateLimit unchanged. This field uses software P-state numbering. See MSRC001_0071[CurPstateLimit] and 2.5.3.1.2.1 [Software P-state Numbering] .

MSRC001_0062 P-state Control

Bits	Description
63:3	MBZ.
2:0	PstateCmd: P-state change command. Read-write; Not-same-for-all . Cold reset value varies by product; after a warm reset, value initializes to the P-state the core was in prior to the reset. Writes to this field cause the core to change to the indicated non-boosted P-state number, specified by MSRC001_00[6B:64] . 0=P0, 1=P1, etc. P-state limits are applied to any P-state requests made through this register. Reads from this field return the last written value, regardless of whether any limits are applied. This field uses software P-state numbering. See 2.5.2.1 [Dependencies Between Cores] , 2.5.3 [CPU Power Management] and 2.5.3.1.2.1 [Software P-state Numbering] .

MSRC001_0063 P-state Status

[SharedC](#).

Bits	Description
63:3	RAZ.
2:0	CurPstate: current P-state. Read; GP-write; Updated-by-hardware. Cold reset: Varies by product. This field provides the frequency component of the current non-boosted P-state of the core (regardless of the source of the P-state change, including MSRC001_0062[PstateCmd] ; see 2.5.3.1.6 [Core P-state Transition Behavior] for information on how these interact). 0=P0, 1=P1, etc. The value of this field is updated when the COF transitions to a new value associated with a P-state. This field uses software P-state numbering. See 2.5.3 [CPU Power Management] and 2.5.3.1.2.1 [Software P-state Numbering] .

MSRC001_00[6B:64] P-state [7:0]

Per-node. Cold reset: Product-specific. Each of these registers specify the frequency and voltage associated with each of the core P-states.

Table 240: Register Mapping for MSRC001_00[6B:64]

Register	Function
MSRC001_0064	P-state 0
MSRC001_0065	P-state 1
MSRC001_0066	P-state 2
MSRC001_0067	P-state 3
MSRC001_0068	P-state 4
MSRC001_0069	P-state 5
MSRC001_006A	P-state 6
MSRC001_006B	P-state 7

The CpuVid field in these registers is required to be programmed to the same value in all cores of a processor, but are allowed to be different between processors in a multi-processor system. All other fields in these registers are required to be programmed to the same value in each core of the coherent fabric. See [2.5.3 \[CPU Power Management\]](#).

When [D18F4x15C\[BoostLock\]=1](#), [MSRC001_00\[6B:64\]\[CpuVid, CpuDid, CpuFid\]](#) have special write requirements associated with them.

Table 241: P-state Definitions

Term	Definition
CoreCOF	Core current operating frequency in MHz. $\text{CoreCOF} = 100 * (\text{MSRC001_00[6B:64][CpuFid]} + 10h) / (2^{\text{MSRC001_00[6B:64][CpuDid]})$.

Bits	Description										
63	PstateEn. Read-write. 1=The P-state specified by this MSR is valid. 0=The P-state specified by this MSR is not valid. The purpose of this register is to indicate if the rest of the P-state information in the register is valid after a reset; it controls no hardware.										
62:42	RAZ.										
41:40	IddDiv: current divisor. Read-write. See IddValue.										
39:32	IddValue: current value. Read-write. After a reset, IddDiv and IddValue combine to specify the expected maximum current dissipation of a single core that is in the P-state corresponding to the MSR number. These values are intended to be used to create ACPI-defined <code>_PSS</code> objects (see 2.5.3.1.9.3 [ACPI Processor P-state Objects]) and to perform the 2.5.3.1.8 [Processor-Systemboard Power Delivery Compatibility Check] . The values are expressed in amps; they are not intended to convey final product power levels. These fields may be subsequently altered by software; they do not affect the hardware behavior. These fields are encoded as follows: <table border="1" data-bbox="267 1764 909 1944"> <thead> <tr> <th><u>IddDiv</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>IddValue / 1 A, Range: 0 to 255 A.</td> </tr> <tr> <td>01b</td> <td>IddValue / 10 A, Range: 0 to 25.5 A.</td> </tr> <tr> <td>10b</td> <td>IddValue / 100 A, Range: 0 to 2.55 A.</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>IddDiv</u>	<u>Description</u>	00b	IddValue / 1 A, Range: 0 to 255 A.	01b	IddValue / 10 A, Range: 0 to 25.5 A.	10b	IddValue / 100 A, Range: 0 to 2.55 A.	11b	Reserved
<u>IddDiv</u>	<u>Description</u>										
00b	IddValue / 1 A, Range: 0 to 255 A.										
01b	IddValue / 10 A, Range: 0 to 25.5 A.										
10b	IddValue / 100 A, Range: 0 to 2.55 A.										
11b	Reserved										

31:23	RAZ.														
22	NbPstate: Northbridge P-state. IF (MSRC001_0071[NbPstateDis]) THEN Read-only. ELSE Read-write. ENDIF. 1=Low performance NB P-state. 0=High performance NB P-state. If this bit is set in any given P-state register, then it must also be set in all enabled lower performance P-state registers as well. Equivalent P-states in each core must program this bit to the same value. See 2.5.4.1 [NB P-states] and D18F5x170[NbPstateThreshold, NbPstateLo, NbPstateHi].														
21:17	RAZ.														
16	CpuVid[7]: core VID bit[7]. Read-write. See CpuVid[6:0].														
15:9	CpuVid[6:0]: core VID. Read-write. See 2.5.1 [Processor Power Planes And Voltage Control]. Writes outside the D18F5x17C[MaxVid, MinVid] range are ignored. If D18F4x15C[BoostLock]=1, then CpuVid for boosted P-states can only be written with values that are greater than (lower voltages than) or equal to the reset value in CpuVid.														
8:6	CpuDid: core divisor ID. Read-write. Specifies the core frequency divisor; see CpuFid. Writes of reserved values are ignored. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Divide-by 1</td> </tr> <tr> <td>1h</td> <td>Divide-by 2</td> </tr> <tr> <td>2h</td> <td>Divide-by 4</td> </tr> <tr> <td>3h</td> <td>Divide-by 8</td> </tr> <tr> <td>4h</td> <td>Divide-by 16</td> </tr> <tr> <td>7h-5h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	Divide-by 1	1h	Divide-by 2	2h	Divide-by 4	3h	Divide-by 8	4h	Divide-by 16	7h-5h	Reserved
Bits	Description														
0h	Divide-by 1														
1h	Divide-by 2														
2h	Divide-by 4														
3h	Divide-by 8														
4h	Divide-by 16														
7h-5h	Reserved														
5:0	CpuFid: core frequency ID. Read-write. Specifies the core frequency multiplier. The core COF is a function of CpuFid and CpuDid, and defined by CoreCOF. This field and CpuDid must be programmed to the requirements specified in MSRC001_0071[MaxCpuCof] and D18F3xD4[MaxSwPstateCpuCof]. Writes to a non-boosted P-state with frequencies greater than D18F3xD4[MaxSwPstateCpuCof] are ignored. Writes to a boosted P-state with frequencies greater than MSRC001_0071[MaxCpuCof] are ignored. If D18F4x15C[BoostLock]=1, then CpuDid for boosted P-states can only be written with values that are greater than or equal to the reset value in CpuDid. If D18F4x15C[BoostLock]=1, then CpuFid for boosted P-states can only be written with values that are less than or equal to the reset value in CpuFid. See 2.5.5 [P-state Bandwidth Requirements].														

MSRC001_0070 COFVID Control

Cold reset: Product-specific. There is one register implemented for each core. This register includes several fields that are identical to MSRC001_00[6B:64]. It is controlled by hardware for P-state transitions. It may also be used by software to directly control the current COF or VID. Accesses to this register that result in invalid COFs or VIDs are ignored. See 2.5.3 [CPU Power Management].

Bits	Description
63:32	RAZ.
31:24	NbVid: Northbridge VID. IF (MSRC001_0071[NbPstateDis]) THEN Read-only. ELSE Read-write. ENDIF. See D18F5x1[6C:60][NbVid].
23	RAZ.
22	NbPstate: Northbridge P-state. IF (MSRC001_0071[NbPstateDis]) THEN Read-only. ELSE Read-write. ENDIF. See MSRC001_00[6B:64][NbPstate].
21	RAZ.

20	CpuVid[7] . Read-write. See CpuVid[6:0].
19	RAZ.
18:16	PstateId: P-state identifier . Read-write. This field is required to provide the P-state number that is associated with the values of the other fields in this register. This value is used by the logic to determine if the P-state is increasing or decreasing. This field uses hardware P-state numbering. See 2.5.3.1.2.2 [Hardware P-state Numbering] .
15:9	CpuVid[6:0]: core VID . Read-write. See MSRC001_00[6B:64][CpuVid] . CpuVid[7:0] = {CpuVid[7], CpuVid[6:0]}.
8:6	CpuDid: core divisor ID . Read-write. See MSRC001_00[6B:64][CpuDid] . The PstateId field must be updated to cause a new CpuDid value to take effect.
5:0	CpuFid: core frequency ID . Read-write. See MSRC001_00[6B:64][CpuFid] . The PstateId field must be updated to cause a new CpuFid value to take effect.

MSRC001_0071 COFVID Status

Read-only. See [2.5.3 \[CPU Power Management\]](#).

Bits	Description
63:59	MaxNbCof: maximum NB COF . Cold reset: Product-specific. Specifies the maximum NB COF supported by the processor. If MaxNbCof is greater than zero, the maximum frequency is 100 MHz * MaxNbCof; if MaxNbCof = 00h, then there is no frequency limit. Any attempt to change the NB COF to a frequency greater than specified by this field is ignored.
58:56	CurPstateLimit: current P-state limit . Updated-by-hardware. Provides the current lowest-performance P-state limit number. This register uses hardware P-state numbering. See MSRC001_0061[CurPstateLimit] and 2.5.3.1.2.2 [Hardware P-state Numbering] .
55	Reserved.
54:49	MaxCpuCof: maximum core COF . Cold reset: Product-specific. Specifies the maximum CPU COF supported by the processor. The maximum frequency is 100 MHz * MaxCpuCof, if MaxCpuCof is greater than zero; if MaxCpuCof = 00h, then there is no frequency limit. Any attempt to change a CPU COF to a frequency greater than specified by this field is ignored.
48:35	Reserved.
34:32	StartupPstate: startup P-state number . Cold reset: Product-specific. Specifies the cold reset VID, FID and DID for the core based on the P-state number selected. StartupPstate uses hardware P-state numbering. See MSRC001_00[6B:64] and 2.5.3.1.2.2 [Hardware P-state Numbering] .
31:24	CurNbVid[7:0]: current NB VID . Updated-by-hardware. Cold reset: Product-specific. This field specifies the current VDDNB voltage. MSRC001_0071[CurNbVid[6:0]] is an alias of D18F5x174[CurNbVid[6:0]] . MSRC001_0071[CurNbVid[7]] is an alias of D18F5x174[CurNbVid[7]] . .
23	NbPstateDis: NB P-states disabled . Value: D18F5x174[NbPstateDis] . MSRC001_0071[NbPstateDis] is an alias of D18F5x174[NbPstateDis] . 0=NB P-state frequency and voltage changes are supported. See D18F5x170[SwNbPstateLoDis, NbPstateDisOnP0] . 1=NB P-state frequency and voltage changes are disabled.
22:21	Reserved.
20	CurCpuVid[7] . Updated-by-hardware. Cold reset: Product-specific. See CurCpuVid[6:0].
19	Reserved.

18:16	CurPstate: current P-state. Updated-by-hardware. Cold reset: Product-specific. Specifies the current P-state requested by the core. This field uses hardware P-state numbering. See MSRC001_0063[CurPstate] and 2.5.3.1.2.2 [Hardware P-state Numbering] . When a P-state change is requested, the value in this field is updated once all required frequency and voltage transitions are completed.
15:9	CurCpuVid[6:0]: current core VID. Updated-by-hardware. Cold reset: Product-specific. CurCpuVid = {CurCpuVid[7], CurCpuVid[6:0]}. This field specifies the current VDD voltage.
8:6	CurCpuDid: current core divisor ID. Updated-by-hardware. Cold reset: Product-specific. Specifies the current CpuDid of the core. See MSRC001_00[6B:64] . When a P-state change is requested, the value in this field is updated once all required frequency and voltage transitions are completed.
5:0	CurCpuFid: current core frequency ID. Updated-by-hardware. Cold reset: Product-specific. Specifies the current CpuFid of the core. See MSRC001_00[6B:64] . When a P-state change is requested, the value in this field is updated once all required frequency and voltage transitions are completed.

MSRC001_0073 C-state Base Address

Reset: 0000_0000_0000_0000h.

Bits	Description
63:16	Reserved.
15:0	CstateAddr: C-state address. Read-write. Specifies the IO addresses trapped by the core for C-state entry requests. A value of 0 in this field specifies that the core does not trap any IO addresses for C-state entry. Writing values greater than FFF8h into this field result in undefined behavior. All other values cause the core to trap IO addresses CstateAddr through CstateAddr+7. See 2.5.3.2.2 [C-state Request Interface] , D18F4x11[C:8] , and D18F4x11C .

MSRC001_0111 SMM Base Address (SMM_BASE)

Reset: 0000_0000_0003_0000h.

This holds the base of the SMM memory region. The value of this register is stored in the save state on entry into SMM (see [2.4.8.2.5 \[SMM Save State\]](#)) and it is restored on returning from SMM. The 16-bit CS (code segment) selector is loaded with SmmBase[19:4] on entering SMM. SmmBase[3:0] is required to be 0. The SMM base address can be changed in two ways:

- The SMM base address, at offset FF00h in the SMM state save area, may be changed by the SMI handler. The RSM instruction updates SmmBase with the new value.
- Normal WRMSR access to this register.

Bits	Description
63:32	Reserved.
31:0	SmmBase. IF MSRC001_0015[SmmLock] THEN Read-only ELSE Read-write ENDIF.

MSRC001_0112 SMM TSeg Base Address (SMMAddr)

Reset: 0000_0000_0000_0000h.

See [2.4.8.2 \[System Management Mode \(SMM\)\]](#) and [2.4.5.1 \[Memory Access to the Physical Address Space\]](#). See [MSRC001_0113](#) for more information about the ASeg and TSeg address ranges.

Each CPU access, directed at CPUAddr, is determined to be in the TSeg range if the following is true:

$\text{CPUAddr}[47:17] \& \text{TSegMask}[47:17] == \text{TSegBase}[47:17] \& \text{TSegMask}[47:17]$.

For example, if TSeg spans 256 KB and starts at the 1 MB address. The [MSRC001_0112](#)[TSegBase] would be set to 0010_0000h and the [MSRC001_0113](#)[TSegMask] to FFFC_0000h (with zeros filling in for bits[16:0]). This results in a TSeg range from 0010_0000 to 0013_FFFFh.

Bits	Description
63:48	Reserved.
47:17	TSegBase[47:17]: TSeg address range base. IF MSRC001_0015 [SmmLock] THEN Read-only ELSE Read-write ENDIF.
16:0	Reserved.

MSRC001_0113 SMM TSeg Mask (SMMMask)

Reset: 0000_0000_0000_0000h.

See [2.4.8.2 \[System Management Mode \(SMM\)\]](#).

The ASeg address range is located at a fixed address from A0000h–BFFFFh. The TSeg range is located at a variable base (specified by [MSRC001_0112](#)[TSegBase]) with a variable size (specified by [MSRC001_0113](#)[TSegMask]). These ranges provide a safe location for SMM code and data that is not readily accessible by non-SMM applications. The SMI handler can be located in one of these two ranges, or it can be located outside these ranges. These ranges must never overlap each other.

This register specifies how accesses to the ASeg and TSeg address ranges are controlled as follows:

- If [A, T]Valid=1, then:
 - If in SMM, then:
 - If [A, T]Close=0, then the accesses are directed to DRAM with memory type as specified in [A, T]MTypeDram.
 - If [A, T]Close=1, then instruction accesses are directed to DRAM with memory type as specified in [A, T]MTypeDram and data accesses are directed at MMIO space and with attributes based on [A, T]MTypeIoWc.
 - If not in SMM, then the accesses are directed at MMIO space with attributes based on [A, T]MTypeIoWc.
- See [2.4.5.1.1 \[Determining Memory Type\]](#).

Bits	Description
63:48	Reserved.
47:17	TSegMask[47:17]: TSeg address range mask. IF MSRC001_0015 [SmmLock] THEN Read-only ELSE Read-write ENDIF. See MSRC001_0112 .
16:15	Reserved.
14:12	TMTypeDram: TSeg address range memory type. IF MSRC001_0015 [SmmLock] THEN Read-only. ELSE Read-write. ENDIF. Specifies the memory type for SMM accesses to the TSeg range that are directed to DRAM. See: Table 208 [Memory Type Definition] .
11	Reserved.

10:8	AMTypeDram: ASeg Range Memory Type. IF MSRC001_0015 [SmmLock] THEN Read-only. ELSE Read-write. ENDIF. Specifies the memory type for SMM accesses to the ASeg range that are directed to DRAM. See: Table 208 [Memory Type Definition] .
7:6	Reserved.
5	TMTypeIoWc: non-SMM TSeg address range memory type. IF MSRC001_0015 [SmmLock] THEN Read-only. ELSE Read-write. ENDIF. Specifies the attribute of TSeg accesses that are directed to MMIO space. 0=UC (uncacheable). 1=WC (write combining).
4	AMTypeIoWc: non-SMM ASeg address range memory type. IF MSRC001_0015 [SmmLock] THEN Read-only. ELSE Read-write. ENDIF. Specifies the attribute of ASeg accesses that are directed to MMIO space. 0=UC (uncacheable). 1=WC (write combining).
3	TClose: send TSeg address range data accesses to MMIO. Read-write. 1=When in SMM, direct data accesses in the TSeg address range to MMIO space. See AClose.
2	AClose: send ASeg address range data accesses to MMIO. Read-write. 1=When in SMM, direct data accesses in the ASeg address range to MMIO space. [A, T]Close allows the SMI handler to access the MMIO space located in the same address region as the [A, T]Seg. When the SMI handler is finished accessing the MMIO space, it must clear the bit. Failure to do so before resuming from SMM causes the CPU to erroneously read the save state from MMIO space.
1	TValid: enable TSeg SMM address range. IF MSRC001_0015 [SmmLock] THEN Read-only. ELSE Read-write. ENDIF. 1=The TSeg address range SMM enabled.
0	AValid: enable ASeg SMM address range. IF MSRC001_0015 [SmmLock] THEN Read-only. ELSE Read-write. ENDIF. 1=The ASeg address range SMM enabled.

MSRC001_0114 Virtual Machine Control (VM_CR)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:32	Reserved.
31:5	MBZ.
4	SvmeDisable: SVM disable. 1= MSRC000_0080 [SVME] must be zero when writing to MSRC000_0080 . Setting this bit when MSRC000_0080 [SVME]=1 causes a #GP fault, regardless of the state of Lock. 0= MSRC000_0080 [SVME] is read-write. See Lock.
3	Lock: SVM lock. Read; write-1-only; cleared-by-hardware. See MSRC001_0118 . 1=SvmeDisable is read-only. 0=SvmeDisable is read-write.
2	DisA20m: disable A20 masking. Read-write; set-by-hardware. 1=Disables A20 masking. This bit is set by hardware when the SKINIT instruction is executed.
1	InterceptInit: intercept INIT. Read-write; set-by-hardware. This bit controls how INIT is delivered in host mode. This bit is set by hardware when the SKINIT instruction is executed. 0=INIT delivered normally. 1=INIT translated into a SX interrupt.
0	Dpd: debug port disable. Read-write; set-by-hardware. Set by hardware when the SKINIT instruction is executed. 1=HDT is disabled. 0=HDT may be enabled.

MSRC001_0115 IGNNE

Bits	Description
63:32	Reserved.
31:1	MBZ.
0	IGNNE: current IGNNE state. Read-write. Reset: X. This bit controls the current state of the processor internal IGNNE signal.

MSRC001_0116 SMM Control (SMM_CTL)

IF (MSRC001_0015[SmmLock]) THEN GP-read-write. ELSE GP-read; write-only. ENDIF.

The bits in this register are processed in the order of: SmmEnter, SmiCycle, SmmDismiss, RsmCycle and SmmExit. However, only the following combination of bits may be set in a single write (all other combinations result in undefined behavior):

- SmmEnter and SmiCycle.
- SmmEnter and SmmDismiss.
- SmmEnter, SmiCycle and SmmDismiss.
- SmmExit and RsmCycle.

Software is responsible for ensuring that SmmEnter and SmmExit operations are properly matched and are not nested.

Bits	Description
63:5	MBZ.
4	RsmCycle: send RSM special cycle. 1=Send a RSM special cycle.
3	SmmExit: exit SMM. 1=Exit SMM.
2	SmiCycle: send SMI special cycle. 1=Send a SMI special cycle.
1	SmmEnter: enter SMM. 1=Enter SMM.
0	SmmDismiss: clear SMI. 1=Clear the SMI pending flag.

MSRC001_0117 Virtual Machine Host Save Physical Address (VM_HSAVE_PA)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:48	MBZ.
47:12	VM_HSAVE_PA: physical address of host save area. Read-write. This register contains the physical address of a 4-KB region where VMRUN saves host state and where vm-exit restores host state from. Writing this register causes a #GP if (FF_FFFF_Fh >= VM_HSAVE_PA >= FD_0000_0h).
11:0	MBZ.

MSRC001_0118 SVM Lock Key

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	SvmLockKey: SVM lock key. RAZ; write. Writes to this register when MSRC001_0114[Lock]=0 write the SvmLockKey. Writes to this register when MSRC001_0114[Lock]=1 and SvmLockKey!=0 cause hardware to clear MSRC001_0114[Lock] if the value written is the same as the value stored in SvmLockKey.

MSRC001_011A Local SMI Status

Reset: 0000_0000_0000_0000h. This registers returns the same information that is returned in [SMMFEC4 \[Local SMI Status\]](#) portion of the SMM save state. The information in this register is only updated when [MSRC001_0116\[SmmDismiss\]](#) is set by software.

Bits	Description
63:32	Reserved.
31:0	See SMMFEC4 [Local SMI Status] .

MSRC001_0140 OS Visible Work-around MSR0 (OSVW_ID_Length)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:16	Reserved.
15:0	OSVWidLength: OS visible work-around ID length. Read-write. See the <i>Revision Guide for AMD Family 15h Models 10h-1Fh Processors</i> for the definition of this field.

MSRC001_0141 OS Visible Work-around MSR1 (OSVW Status)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	OsvwStatusBits: OS visible work-around status bits. Read-write. See the <i>Revision Guide for AMD Family 15h Models 10h-1Fh Processors</i> for the definition of this field.

MSRC001_020[A,8,6,4,2,0] Performance Event Select (PERF_CTL[5:0])

Reset: 0000_0000_0000_0000h. See [2.6.1 \[Core Performance Monitor Counters\]](#). [MSRC001_00\[03:00\]](#) is an alias of [MSRC001_020\[6,4,2,0\]](#). To accurately start counting with the write that enables the counter, disable the counter when changing the event and then enable the counter with a second MSR write.

Table 242: Register Mapping for [MSRC001_020\[A,8,6,4,2,0\]](#)

Register	Function
MSRC001_0200	Counter 0
MSRC001_0202	Counter 1
MSRC001_0204	Counter 2
MSRC001_0206	Counter 3
MSRC001_0208	Counter 4
MSRC001_020A	Counter 5

The edge count mode increments the counter when a transition happens on the monitored event. If the event selected is changed without disabling the counter, an extra edge is falsely detected when the first event is a static 0 and the second event is a static one. To avoid this false edge detection, disable the counter when changing the event and then enable the counter with a second MSR write.

Bits	Description										
63:42	Reserved.										
41:40	HostGuestOnly: count only host/guest events. Read-write. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Count all events, irrespective of guest/host.</td> </tr> <tr> <td>01b</td> <td>Count guest events if MSRC000_0080[SVME]=1.</td> </tr> <tr> <td>10b</td> <td>Count host events if MSRC000_0080[SVME]=1.</td> </tr> <tr> <td>11b</td> <td>Count all guest and host events if MSRC000_0080[SVME]=1.</td> </tr> </tbody> </table>	Bits	Description	00b	Count all events, irrespective of guest/host.	01b	Count guest events if MSRC000_0080[SVME] =1.	10b	Count host events if MSRC000_0080[SVME] =1.	11b	Count all guest and host events if MSRC000_0080[SVME] =1.
Bits	Description										
00b	Count all events, irrespective of guest/host.										
01b	Count guest events if MSRC000_0080[SVME] =1.										
10b	Count host events if MSRC000_0080[SVME] =1.										
11b	Count all guest and host events if MSRC000_0080[SVME] =1.										
39:36	Reserved.										
35:32	EventSelect[11:8]: performance event select. See: EventSelect[7:0].										
31:24	CntMask: counter mask. Read-write. Controls the number of events counted per clock cycle. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>The corresponding PERF_CTR[5:0] register increments by the number of events occurring in a clock cycle. Maximum number of events in one cycle is 32.</td> </tr> <tr> <td>1Fh-01h</td> <td>When Inv = 0, the corresponding PERF_CTR[5:0] register increments by 1, if the number of events occurring in a clock cycle is greater than or equal to the CntMask value. When Inv = 1, the corresponding PERF_CTR[5:0] register increments by 1, if the number of events occurring in a clock cycle is less than CntMask value.</td> </tr> <tr> <td>FFh-20h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00h	The corresponding PERF_CTR[5:0] register increments by the number of events occurring in a clock cycle. Maximum number of events in one cycle is 32.	1Fh-01h	When Inv = 0, the corresponding PERF_CTR[5:0] register increments by 1, if the number of events occurring in a clock cycle is greater than or equal to the CntMask value. When Inv = 1, the corresponding PERF_CTR[5:0] register increments by 1, if the number of events occurring in a clock cycle is less than CntMask value.	FFh-20h	Reserved		
Bits	Description										
00h	The corresponding PERF_CTR[5:0] register increments by the number of events occurring in a clock cycle. Maximum number of events in one cycle is 32.										
1Fh-01h	When Inv = 0, the corresponding PERF_CTR[5:0] register increments by 1, if the number of events occurring in a clock cycle is greater than or equal to the CntMask value. When Inv = 1, the corresponding PERF_CTR[5:0] register increments by 1, if the number of events occurring in a clock cycle is less than CntMask value.										
FFh-20h	Reserved										
23	Inv: invert counter mask. Read-write. See CntMask.										
22	En: enable performance counter. Read-write. 1= Performance event counter is enabled.										
21	Reserved.										
20	Int: enable APIC interrupt. Read-write. 1=APIC performance counter LVT interrupt is enabled to generate an interrupt via APIC340 when the performance counter overflows.										
19	MBZ.										
18	Edge: edge detect. Read-write. 0=Level detect. 1=Edge detect.										
17:16	OsUserMode: OS and user mode. Read-write. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Count no events.</td> </tr> <tr> <td>01b</td> <td>Count user events (CPL>0).</td> </tr> <tr> <td>10b</td> <td>Count OS events (CPL=0).</td> </tr> <tr> <td>11b</td> <td>Count all events, irrespective of the CPL.</td> </tr> </tbody> </table>	Bits	Description	00b	Count no events.	01b	Count user events (CPL>0).	10b	Count OS events (CPL=0).	11b	Count all events, irrespective of the CPL.
Bits	Description										
00b	Count no events.										
01b	Count user events (CPL>0).										
10b	Count OS events (CPL=0).										
11b	Count all events, irrespective of the CPL.										

15:8	UnitMask: event qualification. Read-write. Each UnitMask bit further specifies or qualifies the event specified by EventSelect. All events selected by UnitMask are simultaneously monitored. Unless otherwise stated, the UnitMask values shown may be combined (logically ORed) to select any desired combination of the sub-events for a given event. In some cases, certain combinations can result in misleading counts, or the UnitMask value is an ordinal rather than a bit mask. These situations are described where applicable, or should be obvious from the event descriptions. For events where no UnitMask table is shown, the UnitMask is Unused .
7:0	EventSelect[7:0]: event select. Read-write. EventSelect[11:0] = {EventSelect[11:8], EventSelect[7:0]}. EventSelect specifies the event or event duration in a processor unit to be counted by the corresponding PERF_CTR[5:0] register. The events are specified in 3.22 [Core Performance Counter Events] . Some events are reserved; when a reserved event is selected, the results are undefined.

MSRC001_020[B,9,7,5,3,1] Performance Event Counter (PERF_CTR[5:0])

Reset: 0000_0000_0000_0000h. See [MSRC001_020\[A,8,6,4,2,0\] \[Performance Event Select \(PERF_CTL\[5:0\]\)\]](#). [MSRC001_00\[07:04\]](#) is an alias of [MSRC001_020\[7,5,3,1\]](#).

Table 243: Register Mapping for [MSRC001_020\[B,9,7,5,3,1\]](#)

Register	Function
MSRC001_0201	Counter 0
MSRC001_0203	Counter 1
MSRC001_0205	Counter 2
MSRC001_0207	Counter 3
MSRC001_0209	Counter 4
MSRC001_020B	Counter 5

Bits	Description
63:48	RAZ.
47:0	CTR: performance counter value. Read-write.

MSRC001_024[6,4,2,0] Northbridge Performance Event Select (NB_PERF_CTL[3:0])

Reset: 0000_0000_0000_0000h. [MSRC001_024\[6,4,2,0\]\[31:0\]](#) is an alias of [D18F5x\[70,60,50,40\]](#). [MSRC001_024\[6,4,2,0\]\[63:32\]](#) is an alias of [D18F5x\[74,64,54,44\]](#).

Table 244: Register Mapping for [MSRC001_024\[6,4,2,0\]](#)

Register	Function
MSRC001_0240	Counter 0
MSRC001_0242	Counter 1
MSRC001_0244	Counter 2
MSRC001_0246	Counter 3

These registers are used to specify the events counted by the [MSRC001_024\[7,5,3,1\] \[Northbridge Performance Event Counter \(NB_PERF_CTR\[3:0\]\)\]](#) and to control other aspects of their operation. Each performance counter supported has a corresponding event-select register that controls its operation. [3.23 \[NB](#)

[Performance Counter Events](#)] shows the events and unit masks supported by the processor.

To accurately start counting with the write that enables the counter, disable the counter when changing the event and then enable the counter with a second MSR write.

The edge count mode increments the counter when a transition happens on the monitored event. If the event selected is changed without disabling the counter, an extra edge is falsely detected when the first event is a static 0 and the second event is a static one. To avoid this false edge detection, disable the counter when changing the event and then enable the counter with a second MSR write.

The northbridge performance counter registers can be used to track events in the northbridge. Northbridge events include all memory controller events and crossbar events as documented in [3.23 \[NB Performance Counter Events\]](#). Since the northbridge performance counter register are shared by all cores, monitoring of northbridge events should only be performed by one core.

Bits	Description
63:41	Reserved.
40:37	IntCoreSel: interrupt to core select. Read-write. Specifies the core to direct the interrupt. Values 0 to N correspond to core 0 to N; The most significant value N is indicated by CpuCoreNum, as defined in section 2.4.3 [Processor Cores and Downcoring] ; All values greater than N are reserved.
36	IntCoreEn: interrupt to core enable. Read-write. 1=Interrupt to a single core specified by IntCoreSel. 0=Interrupt to all cores.
35:32	EventSelect[11:8]: performance event select. Read-write. See: EventSelect[7:0].
31:23	Reserved.
22	En: enable performance counter. Read-write. 1= Performance event counter is enabled.
21	Reserved.
20	Int: enable APIC interrupt. Read-write. 1=APIC performance counter LVT interrupt is enabled to generate an interrupt via APIC340 to all local APIC's on this node when the performance counter overflows.
19	Reserved.
18:16	Reserved.
15:8	UnitMask: event qualification. Read-write. Each UnitMask bit further specifies or qualifies the event specified by EventSelect. All events selected by UnitMask are simultaneously monitored. Unless otherwise stated, the UnitMask values shown may be combined (logically ORed) to select any desired combination of the sub-events for a given event. In some cases, certain combinations can result in misleading counts, or the UnitMask value is an ordinal rather than a bit mask. These situations are described where applicable, or should be obvious from the event descriptions. For events where no UnitMask table is shown, the UnitMask is Unused .
7:0	EventSelect[7:0]: event select. Read-write. This field, along with EventSelect[11:8] above, combine to form the 12-bit event select field, EventSelect[11:0]. EventSelect specifies the event or event duration in a processor unit to be counted by the corresponding NB_PERF_CTR[3:0] register. The events are specified in 3.23 [NB Performance Counter Events] . Some events are reserved; when a reserved event is selected, the results are undefined.

MSRC001_024[7,5,3,1] Northbridge Performance Event Counter (NB_PERF_CTR[3:0])

Reset: 0000_0000_0000_0000h. [MSRC001_024\[7,5,3,1\]\[31:0\]](#) is an alias of [D18F5x\[78,68,58,48\]](#).

MSRC001_024[7,5,3,1][63:32] is an alias of D18F5x[7C,6C,5C,4C].

Table 245: Register Mapping for MSRC001_024[7,5,3,1]

Register	Function
MSRC001_0241	Counter 0
MSRC001_0243	Counter 1
MSRC001_0245	Counter 2
MSRC001_0247	Counter 3

The northbridge provides four 48-bit performance counters. Each counter can monitor a different event specified by MSRC001_024[6,4,2,0] [Northbridge Performance Event Select (NB_PERF_CTL[3:0])]. The accuracy of the counters is not ensured.

Performance counters are used to count specific processor events, such as data-cache misses, or the duration of events, such as the number of clocks it takes to return data from memory after a cache miss. During event counting, the processor increments the counter when it detects an occurrence of the event. During duration measurement, the processor counts the number of processor clocks it takes to complete an event. Each performance counter can be used to count one event, or measure the duration of one event at a time.

In addition to the RDMSR instruction, the NB_PERF_CTR[3:0] registers can be read using a special read performance-monitoring counter instruction, RDPMC.

Writing the performance counters can be useful if there is an intention for software to count a specific number of events, and then trigger an interrupt when that count is reached. An interrupt can be triggered when a performance counter overflows. Software should use the WRMSR instruction to load the count as a two's-complement negative number into the performance counter. This causes the counter to overflow after counting the appropriate number of times.

The performance counters are not assured of producing identical measurements each time they are used to measure a particular instruction sequence, and they should not be used to take measurements of very small instruction sequences. The RDPMC instruction is not serializing, and it can be executed out-of-order with respect to other instructions around it. Even when bound by serializing instructions, the system environment at the time the instruction is executed can cause events to be counted before the counter value is loaded into EDX:EAX.

Bits	Description
63:48	RAZ.
47:0	CTR: performance counter value. Read-write. Returns the current value of the event counter.

3.21 MSRs - MSRC001_1xxx

MSRC001_1002 CPUID Features for CPUID Fn0000_0007_E[B,A]X_x0

Read-write. Reset: {CPUID Fn0000_0007_EAX_x0, CPUID Fn0000_0007_EBX_x0}.

MSRC001_1002[63:32] provides back-door control over values read from CPUID Fn0000_0007_EAX_x0;

MSRC001_1002[31:0] provides back-door control over values read from CPUID Fn0000_0007_EBX_x0.

Bits	Description
63:4	Reserved.
3	BMI1.
2:0	Reserved.

MSRC001_1003 Thermal and Power Management CPUID Features

MSRC001_1003 provides control over values read from CPUID Fn0000_0006_ECX.

Bits	Description
63:32	Reserved.
31:0	FeaturesEcx. Read-write. Reset: CPUID Fn0000_0006_ECX. Provides back-door control over the features reported in CPUID Fn0000_0006_ECX.

MSRC001_1004 CPUID Features (Features)

Read-write. Reset: {CPUID Fn0000_0001_ECX, CPUID Fn0000_0001_EDX}. MSRC001_1004[63:32] provides back-door control over values read from CPUID Fn0000_0001_ECX; MSRC001_1004[31:0] provides back-door control over values read from CPUID Fn0000_0001_EDX.

Bits	Description
63:62	Reserved.
61	F16C.
60	AVX.
59	OSXSAVE. Modifies CPUID Fn0000_0001_ECX[OSXSAVE] only if CR4[OSXSAVE].
58	XSAVE.
57	AES. Modifies CPUID Fn0000_0001_ECX[AES] only if the reset value is 1 .
56	Reserved.
55	POPCNT.
54	Reserved.
53	x2APIC.
52	SSE42.
51	SSE41.
50:46	Reserved.
45	CMPXCHG16B.
44:42	Reserved.

41	SSSE3.
40:36	Reserved.
35	Monitor. Modifies CPUID Fn0000_0001_ECX[Monitor] only if ~MSRC001_0015[MonMwaitDis] .
34	Reserved.
33	PCLMULQDQ. Modifies CPUID Fn0000_0001_ECX[PCLMULQDQ] only if the reset value is 1
32	SSE3.
31:29	Reserved.
28	HTT.
27	Reserved.
26	SSE2.
25	SSE.
24	FXSR.
23	MMX.
22:20	Reserved.
19	CLFSH.
18	Reserved.
17	PSE36.
16	PAT.
15	CMOV.
14	MCA.
13	PGE.
12	MTRR.
11	SysEnterSysExit.
10	Reserved.
9	APIC. Modifies CPUID Fn0000_0001_EDX[APIC] only if MSR0000_001B[ApicEn] .
8	CMPXCHG8B.
7	MCE.
6	PAE.
5	MSR.
4	TSC.
3	PSE.
2	DE.
1	VME.
0	FPU.

MSRC001_1005 Extended CPUID Features (ExtFeatures)

Read-write. Reset: {[CPUID Fn8000_0001_ECX](#), [CPUID Fn8000_0001_EDX](#)}. [MSRC001_1005\[63:32\]](#) provides back-door control over values read from [CPUID Fn8000_0001_ECX](#); [MSRC001_1005\[31:0\]](#) provides

back-door control over values read from [CPUID Fn8000_0001_EDX](#).

Bits	Description
63:57	Reserved.
56	PerfCtrExtNB.
55	PerfCtrExtCore.
54	TopologyExtensions. BIOS: IF (CPUID Fn8000_0001_EBX [PkgType]==0010b) THEN 1 ELSE 0 ENDIF.
53	TBM.
52	Reserved.
51	NodeId. BIOS: 1.
50	Reserved.
49	Reserved.
48	FMA4.
47	LWP.
46	Reserved.
45	WDT.
44	SKINIT.
43	XOP.
42	IBS.
41	OSVW.
40	3DNowPrefetch.
39	MisAlignSse.
38	SSE4A.
37	ABM.
36	AltMovCr8.
35	ExtApicSpace.
34	SVM. Modifies CPUID Fn8000_0001_ECX [SVM] only if D18F3xE8 [SvmCapable].
33	CmpLegacy.
32	LahfSahf.
31	3DNow.
30	3DNowExt.
29	LM.
28	Reserved.
27	RDTSMP.
26	Page1GB.
25	FFXSR.
24	FXSR.
23	MMX.
22	MmxExt.

21	Reserved.
20	NX.
19:18	Reserved.
17	PSE36.
16	PAT.
15	CMOV.
14	MCA.
13	PGE.
12	MTRR.
11	SysCallSysRet.
10	Reserved.
9	APIC.
8	CMPXCHG8B.
7	MCE.
6	PAE.
5	MSR.
4	TSC.
3	PSE.
2	DE.
1	VME.
0	FPU.

MSRC001_100C Node ID

Reset: 0000_0000_0000_0000h.

Bits	Description								
63:32	Reserved.								
31:12	Reserved.								
11:6	BiosScratch. Read-write. This field has no function and may be used by BIOS as a scratch field.								
5:3	<p>NodesPerProcessor. Read-write. BIOS: 000b. Specifies the number of nodes per processor if CPUID Fn8000_0001_ECX[NodeId]=1.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>1 node per processor</td> </tr> <tr> <td>001b</td> <td>2 nodes per processor</td> </tr> <tr> <td>111b-010b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	1 node per processor	001b	2 nodes per processor	111b-010b	Reserved
Bits	Description								
000b	1 node per processor								
001b	2 nodes per processor								
111b-010b	Reserved								
2:0	<p>NodeId. Read-write. BIOS: D18F0x60[NodeId]. Specifies the node ID if CPUID Fn8000_0001_ECX[NodeId]=1. Software is recommended to use CPUID Fn8000_001E_ECX[NodeId] instead of this field.</p>								

MSRC001_1020 Load-Store Configuration (LS_CFG)

Bits	Description
63:29	Reserved.
28	DisSS . Read-write. Reset: 0. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot]. 1=Disable streaming store functionality.
27:0	Reserved.

MSRC001_1021 Instruction Cache Configuration (IC_CFG)

SharedC. Reset: 0000_0000_0000_0000h.

Bits	Description								
63:40	Reserved.								
39	DisLoopPredictor . Read-write. BIOS: 1. 1=Disable loop predictor.								
38:10	Reserved.								
9	DisSpecTlbRld . Read-write. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot]. 1=Disable speculative IC TLB reload request; the request is not made to the TLB walker until the fetch is non-speculative.								
8:5	Reserved.								
4:1	DisIcWayFilter: disable IC way access filter . Read-write. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Enable IC way access filter</td> </tr> <tr> <td>Eh-1h</td> <td>Reserved</td> </tr> <tr> <td>Fh</td> <td>Disable IC way access filter</td> </tr> </tbody> </table>	Bits	Description	0h	Enable IC way access filter	Eh-1h	Reserved	Fh	Disable IC way access filter
Bits	Description								
0h	Enable IC way access filter								
Eh-1h	Reserved								
Fh	Disable IC way access filter								
0	Reserved.								

MSRC001_1022 Data Cache Configuration (DC_CFG)

Bits	Description
63:14	Reserved.
13	DisHwPf . Read-write. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot]. Reset: 0. 1=Disable the DC hardware prefetcher.
12:5	Reserved.
4	DisSpecTlbRld . Read-write. Reset: 0. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot]. 1=Disable speculative TLB reloads.
3:0	Reserved.

MSRC001_1023 Combined Unit Configuration (CU_CFG)

SharedC.

Bits	Description								
63:24	Reserved.								
23	L2WayLock: L2 way lock enable. Read-write. Reset: 0. 1=Allocations and evictions for the L2 ways \geq L2FirstLockedWay are disabled. Probes can still invalidate a line in a locked way. Cache lines in the locked ways of the L2 are still accessible by software. See 2.3.3 [Using L2 Cache as General Storage During Boot].								
22:19	L2FirstLockedWay: first L2 way locked. Read-write. Reset: 0h. See L2WayLock. IF (MSRC001_1023[L2WayLock]) THEN (MSRC001_1023[L2FirstLockedWay] \geq 1). <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved</td> </tr> <tr> <td>Eh-1h</td> <td>Ways \lt L2FirstLockedWay to 15 locked.</td> </tr> <tr> <td>Fh</td> <td>Way 15 locked.</td> </tr> </tbody> </table>	Bits	Description	0h	Reserved	Eh-1h	Ways \lt L2FirstLockedWay to 15 locked.	Fh	Way 15 locked.
Bits	Description								
0h	Reserved								
Eh-1h	Ways \lt L2FirstLockedWay to 15 locked.								
Fh	Way 15 locked.								
18:11	Reserved.								
10	DcacheAggressivePriority. Read-write. Reset: 1. BIOS: 0. 0=I-cache requests are ensured fairness when arbitrating for L2 cache access with respect to D-cache even if the I-cache request is no longer needed due to branch mispredicts or other flushes. 1=D-cache is always aggressively given priority over I-cache requests.								
9:0	Reserved.								

MSRC001_1028 Floating Point Configuration (FP_CFG)

SharedC.

Bits	Description
63:45	Reserved.
44:42	DiDtCfg4. Read-write. Reset: 111b. BIOS: D18F3x1FC[DiDtCfg4].
41	DiDtCfg5. Read-write. Reset: 0. BIOS: D18F3x1FC[DiDtCfg5].
40:35	Reserved.
34:27	DiDtCfg1. Read-write. Reset: 10011011b. BIOS: D18F3x1FC[DiDtCfg1].
26:25	DiDtCfg2. Read-write. Reset: 00b. BIOS: D18F3x1FC[DiDtCfg2].
24:23	Reserved.
22:18	DiDtCfg0. Read-write. Reset: 11111b. BIOS: D18F3x1FC[DiDtCfg0].
17	Reserved.
16	DiDtMode. Read-write. Reset: 0. BIOS: D18F3x1FC[DiDtMode].
15:0	Reserved.

MSRC001_102A Combined Unit Configuration 2 (CU_CFG2)

SharedC.

Bits	Description												
63:51	Reserved.												
50	RdMmExtCfgQwEn: read mmio extended config quadword enable. Read-write. Reset: 0. BIOS: 1. 1=MMIO reads to extended config space do not need to be doubleword aligned and may be up to quadword sized. This is to support 64-bit MMIO reads to extended config space. 0=MMIO reads to extended config space need to be doubleword aligned and may be up to doubleword sized. MMIO reads to extended config space that are either not doubleword aligned or greater than doubleword sized are treated as plain MMIO reads.												
49:38	Reserved.												
37:36	ThrottleNbInterface[3:2]. Read-write. Reset: 01b. BIOS: 00b. See ThrottleNbInterface[1:0].												
35:11	Reserved.												
10	VicResyncChkEn. Read-write. Reset: 0. BIOS: 1. 1=Generate an internal probe to NB for non-shared victims. This bit is required to be set for the Monitor/Mwait instructions.												
9:8	Reserved.												
7:6	ThrottleNbInterface[1:0]. Read-write. Reset: 11b. Specifies how many clocks the CU needs to wait before sending the next packet of information to the NB. This applies to the CU->NB request interface and the CU->NB probe response interface. ThrottleNbInterface[3:0] = {ThrottleNbInterface[3:2], ThrottleNbInterface[1:0]}. BIOS: NumOnes(D18F5x80[Enabled])-1. ThrottleNbInterface[3:0]>=NumOnes(D18F5x80[Enabled])-1. This field must be programmed to a value greater than or equal to the number of compute units in the node that have at least one enabled core minus 1. See 2.4.3 [Processor Cores and Downcoring] . <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0000b</td> <td>0 Clocks.</td> </tr> <tr> <td>0001b</td> <td>1 Clock.</td> </tr> <tr> <td>0010b</td> <td>2 Clocks.</td> </tr> <tr> <td>0011b</td> <td>3 Clocks.</td> </tr> <tr> <td>1111b-0100b</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	0000b	0 Clocks.	0001b	1 Clock.	0010b	2 Clocks.	0011b	3 Clocks.	1111b-0100b	Reserved.
Bits	Description												
0000b	0 Clocks.												
0001b	1 Clock.												
0010b	2 Clocks.												
0011b	3 Clocks.												
1111b-0100b	Reserved.												
5:0	Reserved.												

MSRC001_102B Combined Unit Configuration 3 (CU_CFG3)

SharedC.

Bits	Description
63:50	Reserved.
49	CombineCr0Cd: combine CR0 CD for both cores of a compute unit. Read-write. Reset: 0. BIOS: Must not be set when using L2 cache as general storage; see 2.3.3 [Using L2 Cache as General Storage During Boot] ; must be set before passing control to the OS. 1=In non-SMM mode, the effective CR0.CD for both cores of a compute unit is the OR of the host CR0.CD bits for both cores of a compute unit; In SMM mode, the effective CR0.CD for both cores of a compute unit is the OR of the host and guest CR0.CD bits for both cores of a compute unit.
48:43	Reserved.
42	PwcDisableWalkerSharing. Read-write. Reset: 0. BIOS: 0. 1=Page table walker sharing is disabled. Core 0 uses page walker 0 and Core 1 uses page walker 1.

41:23	Reserved.										
22	PfcDoubleStride. Read-write. Reset: 0. BIOS: 1. 1=Prefetch N and N+1 offsets ahead of a stride miss instead of just N. N is configurable by PfcStrideMul.										
21:20	PfcStrideMul. Read-write. Reset: 01b. Specifies the number of stride offsets that are prefetched.										
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>3</td> </tr> <tr> <td>01b</td> <td>4</td> </tr> <tr> <td>10b</td> <td>5</td> </tr> <tr> <td>11b</td> <td>6</td> </tr> </tbody> </table>	Bits	Description	00b	3	01b	4	10b	5	11b	6
Bits	Description										
00b	3										
01b	4										
10b	5										
11b	6										
19:0	Reserved.										

MSRC001_102D Load-Store Configuration 2 (LS_CFG2)

Reset: 0000_0000_0202_0000h.

Bits	Description
63:24	Reserved.
23	DisScbThreshold. Read-write. BIOS: 1. 1=Disable SCB threshold and force deallocation.
22:15	Reserved.
14	ForceSmcCheckFlwStDis. Read-write. BIOS: D18F3x1FC [ForceSmcCheckFlwStDis]. 0=Force a self modifying code check when a cache probe hits a store that has not retired.
13:0	Reserved.

MSRC001_1030 IBS Fetch Control (IC_IBS_CTL)

Reset: 0000_0000_0000_0000h. See 2.6.3 [Instruction Based Sampling (IBS)].

The IBS fetch sampling engine is described as follows:

- The periodic fetch counter is an internal 20-bit counter:
 - The periodic fetch counter [19:4] is set to IbsFetchCnt[19:4] and the periodic fetch counter [3:0] is set according to IbsRandEn when IbsFetchEn is changed from 0 to 1.
 - It increments for every fetch cycle that completes when IbsFetchEn=1 and IbsFetchVal=0.
 - The periodic fetch counter is undefined when IbsFetchEn=0 or IbsFetchVal=1.
 - When IbsFetchCnt[19:4] is read it returns the current value of the periodic fetch counter [19:4].
- When the periodic fetch counter reaches {IbsFetchMaxCnt[19:4],0h} and the selected instruction fetch completes or is aborted:
 - IbsFetchVal is set to 1.
 - Drivers can't assume that IbsFetchCnt[19:4] is 0 when IbsFetchVal==1.
 - The status of the operation is written to the IBS fetch registers (this register, [MSRC001_1031](#) and [MSRC001_1032](#)).
 - An interrupt is generated as specified by [MSRC001_103A](#). The interrupt service routine associated with this interrupt is responsible for saving the performance information stored in IBS execution registers.

Bits	Description
63:58	Reserved.

57	IbsRandEn: random instruction fetch tagging enable. Read-write. 1=Bits[3:0] of the fetch counter are randomized when IbsFetchEn is set to start the fetch counter. 0=Bits[3:0] of the fetch counter are set to 0h when IbsFetchEn is set to start the fetch counter.										
56	IbsL2TlbMiss: instruction cache L2TLB miss. Read-only. 1=The instruction fetch missed in the L2 TLB.										
55	IbsL1TlbMiss: instruction cache L1TLB miss. Read-only. 1=The instruction fetch missed in the L1 TLB.										
54:53	IbsL1TlbPgSz: instruction cache L1TLB page size. Read-only. Indicates the page size of the translation in the L1 TLB. This field is only valid if IbsPhyAddrValid==1. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>4 KB</td> </tr> <tr> <td>01b</td> <td>2 MB</td> </tr> <tr> <td>10b</td> <td>1 GB</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	4 KB	01b	2 MB	10b	1 GB	11b	Reserved
Bits	Description										
00b	4 KB										
01b	2 MB										
10b	1 GB										
11b	Reserved										
52	IbsPhyAddrValid: instruction fetch physical address valid. Read-only. 1=The physical address in MSRC001_1032 and the IbsL1TlbPgSz field are valid for the instruction fetch.										
51	IbsIcMiss: instruction cache miss. Read-only. 1=The instruction fetch missed in the instruction cache.										
50	IbsFetchComp: instruction fetch complete. Read-only. 1=The instruction fetch completed and the data is available for use by the instruction decoder.										
49	IbsFetchVal: instruction fetch valid. Read-only; updated-by-hardware. 1=New instruction fetch data available. When this bit is set, the fetch counter stops counting and an interrupt is generated as specified by MSRC001_103A . This bit must be cleared for the fetch counter to start counting. To clear this bit, software writes 0000h to IbsFetchCnt[19:4] to start the fetch counter at IbsFetchMaxCnt[19:4].										
48	IbsFetchEn: instruction fetch enable. Read-write. 1=Instruction fetch sampling is enabled.										
47:32	IbsFetchLat: instruction fetch latency. Read-only. Reset: X. Indicates the number of clock cycles from when the instruction fetch was initiated to when the data was delivered to the core. If the instruction fetch is abandoned before the fetch completes, this field returns the number of clock cycles from when the instruction fetch was initiated to when the fetch was abandoned.										
31:16	IbsFetchCnt[19:4]. Read-write; updated-by-hardware. Returns the current value of bits[19:4] of the periodic fetch counter.										
15:0	IbsFetchMaxCnt[19:4]. Read-write. Specifies maximum count value of the periodic fetch counter. Programming this field to 0000h and setting IbsFetchEn results in undefined behavior. Bits 19:4 of the maximum count are programmed in the field. Bits[3:0] of the maximum count are always 0000b.										

MSRC001_1031 IBS Fetch Linear Address (IC_IBS_LIN_AD)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	IbsFetchLinAd: instruction fetch linear address. Read-write; updated-by-hardware. Provides the linear address in canonical form for the tagged instruction fetch.

MSRC001_1032 IBS Fetch Physical Address (IC_IBS_PHYS_AD)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	IbsFetchPhysAd: instruction fetch physical address. Read-write; updated-by-hardware. Provides the physical address for the tagged instruction fetch. The lower 12 bits are not modified by address translation, so they are always the same as the linear address. This field contains valid data only if MSRC001_1030 [IbsPhyAddrValid] is asserted.

MSRC001_1033 IBS Execution Control (SC_IBS_CTL)

Reset: 0000_0000_0000_0000h. See 2.6.3 [Instruction Based Sampling (IBS)].

The IBS execution sampling engine is described as follows for IbsOpCntCtl==1. If IbsOpCntCtl==1n then references to “periodic op counter” mean “periodic cycle counter”.

- The periodic op counter is an internal 27-bit counter:
 - It is set to IbsOpCurCnt[26:0] when IbsOpEn is changed from 0 to 1.
 - It increments every dispatched op when IbsOpEn=1 and IbsOpVal=0.
 - The periodic op counter is undefined when IbsOpEn=0 or IbsOpVal=1.
 - When IbsOpCurCnt[26:0] is read then it returns the current value of the periodic micro-op counter [26:0].
- When the periodic micro-op counter reaches IbsOpMaxCnt:
 - The next dispatched micro-op is tagged if IbsOpCntCtl==1. A valid op in the next dispatched line is tagged if IbsOpCntCtl==0. See IbsOpCntCtl.
 - The periodic micro-op counter [26:7]=0; [6:0] is randomized by hardware.
- The periodic micro-op counter is not modified when a tagged micro-op is flushed.
- When a tagged micro-op is retired:
 - IbsOpVal is set to 1.
 - Drivers can't assume that IbsOpCurCnt is 0 when IbsOpVal==1.
 - The status of the operation is written to the IBS execution registers (this register, [MSRC001_1034](#), [MSRC001_1035](#), [MSRC001_1036](#), [MSRC001_1037](#), [MSRC001_1038](#) and [MSRC001_1039](#)).
 - An interrupt is generated as specified by [MSRC001_103A](#). The interrupt service routine associated with this interrupt is responsible for saving the performance information stored in IBS execution registers.

Bits	Description
63:59	Reserved.
58:32	IbsOpCurCnt[26:0]: periodic op counter current count. Read-write; updated-by-hardware. Returns the current value of the periodic op counter.
31:27	Reserved.
26:20	IbsOpMaxCnt[26:20]: periodic op counter maximum count. See: IbsOpMaxCnt[19:4].
19	IbsOpCntCtl: periodic op counter count control. Read-write. 1=Count dispatched ops; when a roll-over occurs, the counter is preloaded with a pseudorandom 7 bit value between 1 and 127. 0=Count clock cycles; a 1-of-4 round robin counter selects an op in the next dispatch line; if the op pointed to by the round robin counter is invalid, then the next younger valid op is selected.
18	IbsOpVal: micro-op sample valid. Read-write; set-by-hardware. 1=New instruction execution data available; the periodic op counter is disabled from counting. An interrupt may be generated when this bit is set as specified by MSRC001_103A [LvtOffset].
17	IbsOpEn: micro-op sampling enable. Read-write. 1=Instruction execution sampling enabled.

16	Reserved.						
15:0	IbsOpMaxCnt[19:4]: periodic op counter maximum count. Read-write. IbsOpMaxCnt[26:0] = {IbsOpMaxCnt[26:20], IbsOpMaxCnt[19:4], 0000b}. Specifies maximum count value of the periodic op counter. Bits [3:0] of the maximum count are always 0000b. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0008h-0000h</td> <td>Reserved</td> </tr> <tr> <td>FFFFh-0009h</td> <td><IbsOpMaxCnt[19:4]*16> ops</td> </tr> </tbody> </table>	Bits	Description	0008h-0000h	Reserved	FFFFh-0009h	<IbsOpMaxCnt[19:4]*16> ops
Bits	Description						
0008h-0000h	Reserved						
FFFFh-0009h	<IbsOpMaxCnt[19:4]*16> ops						

MSRC001_1034 IBS Op Logical Address (IBSOP_RIP)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	IbsOpRip: micro-op logical address. Read-write; updated-by-hardware. Logical address in canonical form for the instruction that contains the tagged micro-op.

MSRC001_1035 IBS Op Data (SC_IBS_DATA)

Read-write; updated-by-hardware. Reset: 0000_0000_0000_0000h.

Bits	Description
63:39	Reserved.
38	IbsRipInvalid: RIP is invalid. 1=Tagged operation RIP is invalid. Support for this indicated by CPUID Fn8000_001B_EAX[RipInvalidChk] .
37	IbsOpBrnRet: branch micro-op retired. 1=Tagged operation was a branch micro-op that retired.
36	IbsOpBrnMisp: mispredicted branch micro-op. 1=Tagged operation was a branch micro-op that was mispredicted. Qualified by IbsOpBrnRet==1.
35	IbsOpBrnTaken: taken branch micro-op. 1=Tagged operation was a branch micro-op that was taken. Qualified by IbsOpBrnRet==1.
34	IbsOpReturn: return micro-op. 1=Tagged operation was return micro-op. Qualified by IbsOpBrnRet==1.
33:32	Reserved.
31:16	IbsTagToRetCtr: micro-op tag to retire count. This field returns the number of cycles from when the micro-op was tagged to when the micro-op was retired. This field is equal to IbsCompToRetCtr when the tagged micro-op is a NOP.
15:0	IbsCompToRetCtr: micro-op completion to retire count. This field returns the number of cycles from when the micro-op was completed to when the micro-op was retired.

MSRC001_1036 IBS Op Data 2 (IbsOpData2)

Reset: 0000_0000h. Northbridge data is only valid for load operations that miss both the L1 data cache and the L2 cache. If a load operation crosses a cache line boundary, the data returned in this register is the data for the access to the lower cache line.

Bits	Description
63:6	Reserved.

5	NbIbsReqCacheHitSt: IBS cache state. Read-write; updated-by-hardware. Valid when the data source type is Cache(2h). 0=M State. 1=O State.																		
4	NbIbsReqDstNode: IBS request destination node. Read-write; updated-by-hardware. 0=The request is serviced by the NB in the same node as the core. 1=The request is serviced by the NB in a different node than the core. Valid when NbIbsReqSrc is non-zero.																		
3	Reserved.																		
2:0	NbIbsReqSrc: northbridge IBS request data source. Read-write; updated-by-hardware. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>No valid status</td> </tr> <tr> <td>1h</td> <td>Reserved</td> </tr> <tr> <td>2h</td> <td>Cache: data returned from another compute-unit cache on the same node.</td> </tr> <tr> <td>3h</td> <td>DRAM: data returned from DRAM</td> </tr> <tr> <td>4h</td> <td>Reserved for remote cache</td> </tr> <tr> <td>5h</td> <td>Reserved</td> </tr> <tr> <td>6h</td> <td>Reserved</td> </tr> <tr> <td>7h</td> <td>Other: data returned from MMIO/Config/PCI/APIC</td> </tr> </tbody> </table>	Bits	Description	0h	No valid status	1h	Reserved	2h	Cache: data returned from another compute-unit cache on the same node.	3h	DRAM: data returned from DRAM	4h	Reserved for remote cache	5h	Reserved	6h	Reserved	7h	Other: data returned from MMIO/Config/PCI/APIC
Bits	Description																		
0h	No valid status																		
1h	Reserved																		
2h	Cache: data returned from another compute-unit cache on the same node.																		
3h	DRAM: data returned from DRAM																		
4h	Reserved for remote cache																		
5h	Reserved																		
6h	Reserved																		
7h	Other: data returned from MMIO/Config/PCI/APIC																		

MSRC001_1037 IBS Op Data 3 (DC_IBS_DATA, IbsOpData3)

Read-write; updated-by-hardware. Reset: 0000_0000_0000_0000h. If a load or store operation crosses a 128-bit boundary, the data returned in this register is the data for the access to the data below the 128-bit boundary.

Bits	Description
63:48	Reserved.
47:32	IbsDcMissLat: data cache miss latency. Indicates the number of clock cycles from when a miss is detected in the data cache to when the data was delivered to the core. The value returned by this counter is not valid for data cache writes or prefetch instructions.
31:20	Reserved.
19	IbsDcL2tlbHit1G: data cache L2TLB hit in 1G page. 1=The physical address for the tagged load or store operation was present in a 1G page table entry in the data cache L2TLB.
18	IbsDcPhyAddrValid: data cache physical address valid. 1=The physical address in MSRC001_1039 is valid for the load or store operation.
17	IbsDcLinAddrValid: data cache linear address valid. 1=The linear address in MSRC001_1038 is valid for the load or store operation.
16	IbsDcMabHit: MAB hit. 1=The tagged load or store operation hit on an already allocated MAB.
15	IbsDcLockedOp: locked operation. 1=Tagged load or store operation is a locked operation.
14	IbsDcUcMemAcc: UC memory access. 1=Tagged load or store operation accessed uncacheable memory.
13	IbsDcWcMemAcc: WC memory access. Read-write; updated-by-hardware. 1=Tagged load or store operation accessed write combining memory.
12:9	Reserved.
8	IbsDcMisAcc: misaligned access. 1=The tagged load or store operation crosses a 128 bit address boundary.
7	IbsDcMiss: data cache miss. 1=The cache line used by the tagged load or store was not present in the data cache.

6	IbsDcL2TlbHit2M: data cache L2TLB hit in 2M page. 1=The physical address for the tagged load or store operation was present in a 2M page table entry in the data cache L2TLB.
5	IbsDcL1TlbHit1G: data cache L1TLB hit in 1G page. 1=The physical address for the tagged load or store operation was present in a 1G page table entry in the data cache L1TLB.
4	IbsDcL1TlbHit2M: data cache L1TLB hit in 2M page. 1=The physical address for the tagged load or store operation was present in a 2M page table entry in the data cache L1TLB.
3	IbsDcL2TlbMiss: data cache L2TLB miss. 1=The physical address for the tagged load or store operation was not present in the data cache L2TLB.
2	IbsDcL1TlbMiss: data cache L1TLB miss. 1=The physical address for the tagged load or store operation was not present in the data cache L1TLB.
1	IbsStOp: store op. 1=Tagged operation is a store operation.
0	IbsLdOp: load op. 1=Tagged operation is a load operation.

MSRC001_1038 IBS DC Linear Address (DC_IBS_LIN_ADDR)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	IbsDcLinAd. Read-write; updated-by-hardware. Provides the linear address in canonical form for the tagged load or store operation. This field contains valid data only if MSRC001_1037 [IbsDcLinAddrValid] is asserted.

MSRC001_1039 IBS DC Physical Address (DC_IBS_PHYS_ADDR)

Reset: 0000_0000_0000_0000h.

Bits	Description
63:48	RAZ.
47:0	IbsDcPhysAd: load or store physical address. Read-write; updated-by-hardware. Provides the physical address for the tagged load or store operation. The lower 12 bits are not modified by address translation, so they are always the same as the linear address. This field contains valid data only if MSRC001_1037 [IbsDcPhyAddrValid] is asserted.

MSRC001_103A IBS Control

GP-write.

Bits	Description
63:32	Reserved.
31:0	Alias of D18F3x1CC . Read-only.

MSRC001_103B IBS Branch Target Address (BP_IBSTGT_RIP)

Reset: 0000_0000_0000_0000h. Support for this register indicated by [CPUID Fn8000_001B_EAX](#)[BrnTrgt].

Bits	Description
63:0	IbsBrTarget . Read-write; updated-by-hardware. The logical address in canonical form for the branch target. Contains a valid target if non-0. Qualified by MSRC001_1035 [IbsOpBrnRet]==1

3.22 Core Performance Counter Events

This section provides the core performance counter events that may be selected through [MSRC001_020\[A,8,6,4,2,0\] \[Performance Event Select \(PERF_CTL\[5:0\]\)\]\[EventSelect and UnitMask\]](#). See that register and [MSRC001_020\[B,9,7,5,3,1\] \[Performance Event Counter \(PERF_CTR\[5:0\]\)\]](#).

For NB performance counter events see [2.6.2 \[NB Performance Monitor Counters\]](#) and [3.23 \[NB Performance Counter Events\]](#).

3.22.1 PMCx0[1F:00] Events (FP)

PMCx000 FPU Pipe Assignment

[PERF_CTL\[3\]](#). . The number of operations (uops) and dual-pipe uops dispatched to each of the 4 FPU execution pipelines. This event reflects how busy the FPU pipelines are and may be used for workload characterization. This includes all operations performed by x87, MMX, and SSE instructions, including moves. Each increment represents a one-cycle dispatch event. This event is a speculative event. (See [PMCx0CB](#)). Since this event includes non-numeric operations it is not suitable for measuring MFLOPS. The number of events logged per cycle can vary from 0 to 8 and must use [PERF_CTL\[3\]](#).

UnitMask	Description
7	Total number dual-pipe uops assigned to Pipe 3.
6	Total number dual-pipe uops assigned to Pipe 2.
5	Total number dual-pipe uops assigned to Pipe 1.
4	Total number dual-pipe uops assigned to Pipe 0.
3	Total number uops assigned to Pipe 3.
2	Total number uops assigned to Pipe 2.
1	Total number uops assigned to Pipe 1.
0	Total number uops assigned to Pipe 0.

PMCx001 FP Scheduler Empty

[PERF_CTL\[5:3\]](#). This is a speculative event. The number of cycles in which the FPU scheduler is empty. Note that some ops like FP loads bypass the scheduler; see the FP MAS for the full list of “no pipe” ops that bypass the scheduler. Invert this ([MSRC001_020\[A,8,6,4,2,0\]\[Inv\]=1](#)) to count cycles in which at least one FPU operation is present in the FPU.

PMCx003 Retired SSE/BNI Ops

[PERF_CTL\[3\]](#). This is a retire-based event. The number of retired SSE/BNI FLOPS. The number of events logged per cycle can vary from 0 to 32.

UnitMask	Description
7	Double precision multiply-add FLOPS. Multiply-add counts as 2 FLOPS.
6	Double precision divide/square root FLOPS.
5	Double precision multiply FLOPS.
4	Double precision add/subtract FLOPS.

3	Single precision multiply-add FLOPS. Multiply-add counts as 2 FLOPS.
2	Single-precision divide/square root FLOPS.
1	Single-precision multiply FLOPS.
0	Single-precision add/subtract FLOPS.

PMCx004 Number of Move Elimination and Scalar Op Optimization

PERF_CTL[3]. This is a dispatch based speculative event, and is useful for measuring the effectiveness of the Move elimination and Scalar code optimization schemes. The number of events logged per cycle can vary from 0 to 8 and must use **PERF_CTL[3]**.

UnitMask	Description
7:4	Reserved.
3	Number of Scalar ops optimized.
2	Number of Ops that are candidates for optimization (have Z-bit either set or pass).
1	Number of SSE Move Ops eliminated.
0	Number of SSE Move Ops.

PMCx005 Retired Serializing Ops

PERF_CTL[5:3]. The number of serializing ops retired.

UnitMask	Description
7:4	Reserved.
3	x87 control word mispredict traps due to mispredictions in RC or PC, or changes in mask bits.
2	x87 bottom-executing uops retired.
1	SSE control word mispredict traps due to mispredictions in RC, FTZ or DAZ, or changes in mask bits.
0	SSE bottom-executing uops retired.

PMCx006 Number of Cycles that a Bottom-Execute uop is in the FP Scheduler

PERF_CTL[5:3]. This is a speculative event.

3.22.2 PMCx0[3F:20] Events (LS)

PMCx020 Segment Register Loads

PERF_CTL[5:0]. The number of segment register loads performed.

UnitMask	Description
7	Reserved.
6	HS
5	GS
4	FS

3	DS
2	SS
1	CS
0	ES

PMCx021 Pipeline Restart Due to Self-Modifying Code

PERF_CTL[5:0]. The number of pipeline restarts that were caused by self-modifying code (a store that hits any instruction that's been fetched for execution beyond the instruction doing the store).

PMCx022 Pipeline Restart Due to Probe Hit

PERF_CTL[5:0]. The number of pipeline restarts caused by an invalidating probe hitting on a speculative out-of-order load.

PMCx023 Load Queue/Store Queue Full

PERF_CTL[2:0]. The number of cycles that the load queue (LDQ) or store queue (STQ) is full. The load queue holds loads that missed the data cache and are waiting on a refill; the store queue holds stores waiting to retire. This condition stalls further data cache accesses, although such stalls may be overlapped by independent instruction execution.

UnitMask	Description
7:2	Reserved.
1	The number of cycles that the store buffer is full.
0	The number of cycles that the load buffer is full.

PMCx024 Locked Operations

PERF_CTL[5:0]. This event covers locked operations performed and their execution time. The execution time represented by the cycle counts is typically overlapped to a large extent with other instructions. The non-speculative cycles event indicates the time when the locked operation is issued as the oldest request in the queue, and is suitable for event-based profiling of lock operations that tend to miss in the cache.

UnitMask	Description
7:4	Reserved.
3	The number of cycles spent in non-speculative phase, including the cache miss penalty.
2	The number of cycles spent in non-speculative phase, excluding cache miss penalty. In the absence of a cache miss, the UnitMask[3] counter counts 8 cycles more for each locked operation than the UnitMask[2] counter.
1	Reserved.
0	The number of locked instructions executed.

PMCx026 Retired CLFLUSH Instructions

PERF_CTL[5:0]. The number of retired CLFLUSH instructions. This is a non-speculative event.

PMCx027 Retired CPUID Instructions

PERF_CTL[5:0]. The number of CPUID instructions retired.

PMCx029 LS Dispatch

PERF_CTL[5:0]. Counts the number of operations dispatched to the LS unit.

UnitMask	Description
7:3	Reserved.
2	Load-op-Stores.
1	Stores.
0	Loads.

PMCx02A Canceled Store to Load Forward Operations

PERF_CTL[5:0]. Counts the number of canceled store to load forward operations.

UnitMask	Description
7:1	Reserved.
0	Either “store is smaller than load” or “different starting byte but partial overlap”.

PMCx02B SMIs Received

PERF_CTL[5:0]. Counts the number of SMIs received.

PMCx030 Executed CLFLUSH Instructions

PERF_CTL[5:0]. The number of executed CLFLUSH instructions. This is a speculative event.

PMCx032 Misaligned Stores

PERF_CTL[5:0]. The number of misaligned stores.

PMCx034 FP Load Buffer Stall

PERF_CTL[5:0]. The number of loads stalled due to buffer full.

3.22.3 PMCx0[5F:40] Events (DC)**PMCx040 Data Cache Accesses**

PERF_CTL[5:0]. The number of accesses to the data cache for load and store references. This may include certain microcode scratchpad accesses, although these are generally rare. This event is a speculative event. The number of events logged per cycle can vary from 0 to 2.

PMCx041 Data Cache Misses

PERF_CTL[5:0]. The number of data cache references which missed in the data cache. This event is a speculative event. Only the first miss for a given line is included; access attempts by other instructions while the refill is still pending are not included in this event. Each event reflects one 64 B cache line refill, and counts of this event are the same as, or very close to, the combined count for **PMCx042**. The number of events logged per cycle can vary from 0 to 2.

UnitMask	Description
7:2	Reserved.
1	First streaming store to a 64 B cache line.
0	First data cache miss or streaming store to a 64 B cache line.

PMCx042 Data Cache Refills from L2 or System

PERF_CTL[5:0]. The number of data cache refills satisfied from the L2 cache and/or the system. Each increment reflects a 64 B transfer. This event is a speculative event.

UnitMask	Description
7:4	Reserved.
3	Fill with read data error.
2	Reserved.
1	Early valid status turned out to be invalid.
0	Fill with good data. (Final valid status is valid)

PMCx043 Data Cache Refills from System

PERF_CTL[2:0]. The number of L1 cache refills satisfied from the system (system memory or another cache), as opposed to the L2. Each increment reflects a 64 B transfer. This event is a speculative event.

PMCx045 Unified TLB Hit

PERF_CTL[2:0]. The number of TLB accesses that miss in the L1 DTLB or L1 and L2 ITLBs and hit in the unified TLB (UCTLB). This event is a speculative event.

UnitMask	Description
7	Reserved.
6	1 GB unified TLB hit for instruction.
5	2 MB unified TLB hit for instruction.
4	4 KB unified TLB hit for instruction.
3	Reserved.
2	1 GB unified TLB hit for data.
1	2 MB unified TLB hit for data.
0	4 KB unified TLB hit for data.

PMCx046 Unified TLB Miss

PERF_CTL[2:0]. The number of TLB accesses that miss in all TLBs. This event is a speculative event.

UnitMask	Description
7	Reserved.
6	1 GB unified TLB miss for instruction.
5	2 MB unified TLB miss for instruction.
4	4 KB unified TLB miss for instruction.
3	Reserved.
2	1 GB unified TLB miss for data.
1	2 MB unified TLB miss for data.
0	4 KB unified TLB miss for data.

PMCx047 Misaligned Accesses

PERF_CTL[5:0]. The number of data cache accesses that are misaligned. These are accesses which cross an 8 B boundary. They incur an extra cache access (reflected in **PMCx040**), and an extra cycle of latency on reads. This event is a speculative event.

PMCx04B Prefetch Instructions Dispatched

PERF_CTL[5:0]. The number of prefetch instructions dispatched by the decoder. Such instructions may or may not cause a cache line transfer. Any Dcache and L2 accesses, hits and misses by prefetch instructions are included in these types of events. This event is a speculative event.

UnitMask	Description
7:3	Reserved.
2	NTA (PrefetchNTA)
1	Store (PrefetchW)
0	Load (Prefetch, PrefetchT0/T1/T2)

PMCx052 Ineffective Software Prefetches

PERF_CTL[5:0]. The number of software prefetches that did not fetch data outside of the processor core.

UnitMask	Description
7:4	Reserved.
3	Software prefetch hit in the L2.
2:1	Reserved.
0	Software prefetch hit in the L1.

PMCx054 Global TLB Flashes

PERF_CTL[2:0]. This event counts TLB flushes that flush TLB entries that have the global bit set.

3.22.4 PMCx[1,0][7F:60] Events (CU)

PMCx065 Memory Requests by Type

PERF_CTL[2:0]. These events reflect accesses to uncacheable (UC), write-combining (WC), and streaming store (SS) activity to WB memory.

UnitMask	Description
7	Requests to non-cacheable (WC+/SS, but not WC) memory, consisting of reads and 64 B sized buffer flushes.
6:2	Reserved.
1	Requests to non-cacheable (WC, but not WC+/SS) memory, consisting of reads and 64 B sized buffer flushes.
0	Requests to non-cacheable (UC) memory.

PMCx067 Data Prefetcher

PERF_CTL[2:0].

UnitMask	Description
7:4	Reserved.
3:2	Reserved.
1	Prefetch attempts.
0	Reserved.

PMCx068 MAB Requests

PERF_CTL[2:0]. Events **PMCx068** and **PMCx069** reflect utilization of the Miss Address buffers (MABs), which handle IC, DC, TLB, WCC, and WCB related requests. The **UnitMask[BufferID]** is an encoded value which selects one of the MABs. **PMCx068** counts the number of cacheable L2 misses handled by the selected MAB; **PMCx069** counts the number of cycles the selected MAB is busy waiting for the NB response. The average latency seen by the selected MAB is the number of cycles spent waiting (**PMCx069**) divided by the number of requests (**PMCx068**).

UnitMask	Description						
7:0	BufferID. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>27-0</td> <td>MAB ID</td> </tr> <tr> <td>255-28</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	27-0	MAB ID	255-28	Reserved
<u>Bits</u>	<u>Description</u>						
27-0	MAB ID						
255-28	Reserved						

PMCx069 MAB Wait Cycles

PERF_CTL[2:0]. See **PMCx068**.

UnitMask	Description
7:0	BufferID. See: PMCx068[BufferID] .

PMCx06C Response From System on Cache Refills

PERF_CTL[2:0]. The number of responses from the system for cache refill requests. The UnitMask may be used to select specific cache coherency states. Each increment represents one 64 B cache line transferred from the system (DRAM or another cache, including another core on the same node) to the data cache, instruction cache or L2 cache (for data prefetcher and TLB table walks). Modified-state responses may be for Dcache store miss refills, PrefetchW software prefetches, hardware prefetches for a store-miss stream, or Change-to-Dirty requests that get a dirty (Owned) probe hit in another cache. Exclusive responses may be for any Icache refill, Dcache load miss refill, other software prefetches, hardware prefetches for a load-miss stream, or TLB table walks that miss in the L2 cache; Shared responses may be for any of those that hit a clean line in another cache.

UnitMask	Description
7:6	Reserved.
5	Modified unwritten
4	Data Error
3	Owned
2	Shared
1	Modified
0	Exclusive

PMCx06D Outwords Written to System

PERF_CTL[2:0]. The number of OW (16 B) data transfers from the processor to the system. These may be part of a 64 B cache line writeback or a 64 B dirty probe hit response, each of which would cause four increments; or a partial or complete Write Combining buffer flush (Sized Write), which could cause from one to four increments.

UnitMask	Description
7:6	Reserved.
5:1	Reserved.
0	OW write transfer.

PMCx076 CPU Clocks not Halted

PERF_CTL[2:0]. The number of core clocks that the CPU is not in a halted state (due to STPCLK or a HLT instruction). This event allows system idle time to be automatically factored out from IPC (or CPI) measurements, providing the OS halts the CPU when going idle. If the OS goes into an idle loop rather than halting, such calculations are influenced by the IPC of the idle loop. The core clock frequency varies with P-states.

PMCx07D Requests to L2 Cache

PERF_CTL[2:0]. The number of requests to the L2 cache for Icache or Dcache fills, or page table lookups for the TLB. These events reflect only read requests to the L2; writes to the L2 are indicated by **PMCx07E**. See

PMCx081, PMCx082, PMCx083, PMCx041, PMCx042, PMCx043.

UnitMask	Description
7	Reserved.
6	L2 cache prefetcher request
5	Reserved.
4	Canceled request
3	NB probe request
2	TLB fill (page table walks)
1	DC fill
0	IC fill

PMCx07E L2 Cache Misses

PERF_CTL[2:0]. The number of requests that miss in the L2 cache. This may include some amount of speculative activity. The IC-fill-miss and DC-fill-miss events tend to mirror the Icache and Dcache refill-from-system **PMCx083** and **PMCx043**, and tend to include more speculative activity than those events.

UnitMask	Description
7:6	Reserved.
5	Reserved.
4	L2 Cache Prefetcher request
3	Reserved.
2	TLB page table walk
1	DC fill (includes possible replays, whereas PMCx041 does not)
0	IC fill

PMCx07F L2 Fill/Writeback

PERF_CTL[2:0]. Each increment represents a 64 B cache line transfer.

UnitMask	Description
7:3	Reserved.
2	L2 Clean Writebacks to system
1	L2 Writebacks to system (Clean and Dirty)
0	L2 fills from system. Note: Fills for non-temporal software prefetch and WP-memtype fills also are counted in this event even though they don't get cached in L2.

PMCx165 Page Splintering

PERF_CTL[2:0]. Counts the number of TLB reloads where a large page is installed into the TLB as a smaller page size.

UnitMask	Description
7:3	Reserved.

2	Host page size is larger than the guest page size.
1	Splintering due to MTRRs, IORRs, APIC, TOMs or other special address region.
0	Guest page size is larger than the host page size when nested paging is enabled.

PMcX16C L2 Prefetcher Trigger Events

[PERF_CTL\[2:0\]](#).

UnitMask	Description
7:2	Reserved.
1	Store L1 miss seen by prefetcher.
0	Load L1 miss seen by prefetcher.

3.22.5 PMcX[1,0][9F:80] Events (IC)

All instruction cache events are speculative events unless specified otherwise.

PMcX080 Instruction Cache Fetches

[PERF_CTL\[2:0\]](#). The number of instruction cache accesses by the instruction fetcher. Each access is an aligned 32 B read, from which a varying number of instructions may be decoded.

PMcX081 Instruction Cache Misses

[PERF_CTL\[2:0\]](#). The number of instruction fetches and prefetch requests that miss in the instruction cache. This is typically equal to or very close to the sum of events 82h and 83h. Each miss results in a 64 B cache line refill.

PMcX082 Instruction Cache Refills from L2

[PERF_CTL\[2:0\]](#). The number of instruction cache refills satisfied from the L2 cache. Each increment represents one 64 B cache line transfer.

PMcX083 Instruction Cache Refills from System

[PERF_CTL\[2:0\]](#). The number of instruction cache refills from system memory (or another cache). Each increment represents one 64 B cache line transfer.

PMcX084 L1 ITLB Miss, L2 ITLB Hit

[PERF_CTL\[2:0\]](#). The number of instruction fetches that miss in the L1 ITLB but hit in the L2 TLB

PMCx085 L1 ITLB Miss, L2 ITLB Miss

PERF_CTL[2:0]. The number of instruction fetches that miss in both the L1 and L2 TLBs.

UnitMask	Description
7:3	Reserved.
2	Instruction fetches to a 1 GB page.
1	Instruction fetches to a 2 MB page.
0	Instruction fetches to a 4 KB page.

PMCx086 Pipeline Restart Due to Instruction Stream Probe

PERF_CTL[2:0]. The number of pipeline restarts caused by invalidating probes that hit on the instruction stream currently being executed. This would happen if the active instruction stream was being modified by another processor in an MP system - typically a highly unlikely event.

PMCx087 Instruction Fetch Stall

PERF_CTL[2:0]. The number of cycles the instruction fetcher is stalled for the core. This may be for a variety of reasons such as branch predictor updates, unconditional branch bubbles, far jumps and cache misses, instruction fetching for the other core while instruction fetch for this core is stalled, among others. May be overlapped by instruction dispatch stalls or instruction execution, such that these stalls don't necessarily impact performance.

PMCx088 Return Stack Hits

PERF_CTL[2:0]. The number of near return instructions (RET or RET Iw) that get their return address from the return address stack (i.e. where the stack has not gone empty) for the core. This may include cases where the address is incorrect (return mispredicts). This may also include speculatively executed false-path returns. Return mispredicts are typically caused by the return address stack underflowing, however they may also be caused by an imbalance in calls vs. returns, such as doing a call but then popping the return address off the stack.

This event cannot be reliably compared with events C9h and CAh (such as to calculate percentage of return mispredicts due to an empty return address stack), since it may include speculatively executed false-path returns that are not included in those retire-time events.

PMCx089 Return Stack Overflows

PERF_CTL[2:0]. The number of (near) call instructions that cause the return address stack to overflow. When this happens, the oldest entry is discarded. This count may include speculatively executed calls.

PMCx08B Instruction Cache Victims

PERF_CTL[2:0]. The number of cachelines evicted from the instruction cache that cause an L2 write due to changed predecode (start bits); the L2 write due to changed predecode doesn't write the instruction bytes. This event does not count IC evictions with unchanged predecode, which are silently dropped without an L2 write.

PMCx08C Instruction Cache Lines Invalidated

PERF_CTL[2:0]. The number of instruction cache lines invalidated.

UnitMask	Description
7:4	Reserved.
3	SMC invalidating probe that hit on in-flight instructions.
2	SMC invalidating probe that missed on in-flight instructions.
1	Non-SMC invalidating probe that hit on in-flight instructions.
0	Non-SMC invalidating probe that missed on in-flight instructions.

PMCx099 ITLB Reloads

PERF_CTL[2:0]. The number of ITLB reload requests.

PMCx09A ITLB Reloads Aborted

PERF_CTL[2:0]. The number of ITLB reloads aborted.

3.22.6 PMCx0[BF:A0] Events

There are no PMCx0[BF:A0] Events.

3.22.7 PMCx[1,0][DF:C0] Events (EX, DE)**PMCx0C0 Retired Instructions**

PERF_CTL[5:0]. The number of instructions retired (execution completed and architectural state updated). This count includes exceptions and interrupts - each exception or interrupt is counted as one instruction.

PMCx0C1 Retired uops

PERF_CTL[5:0]. The number of micro-ops retired. This includes all processor activity (instructions, exceptions, interrupts, microcode assists, etc.). The number of events logged per cycle can vary from 0 to 4.

PMCx0C2 Retired Branch Instructions

PERF_CTL[5:0]. The number of branch instructions retired. This includes all types of architectural control flow changes, including exceptions and interrupts.

PMCx0C3 Retired Mispredicted Branch Instructions

PERF_CTL[5:0]. The number of branch instructions retired, of any type, that were not correctly predicted. This includes those for which prediction is not attempted (far control transfers, exceptions and interrupts).

PMCx0C4 Retired Taken Branch Instructions

PERF_CTL[5:0]. The number of taken branches that were retired. This includes all types of architectural con-

control flow changes, including exceptions and interrupts.

PMCx0C5 Retired Taken Branch Instructions Mispredicted

PERF_CTL[5:0]. The number of retired taken branch instructions that were mispredicted.

PMCx0C6 Retired Far Control Transfers

PERF_CTL[5:0]. The number of far control transfers retired including far call/jump/return, IRET, SYSCALL and SYSRET, plus exceptions and interrupts. Far control transfers are not subject to branch prediction.

PMCx0C7 Retired Branch Resyncs

PERF_CTL[5:0]. The number of resync branches. These reflect pipeline restarts due to certain microcode assists and events such as writes to the active instruction stream, among other things. Each occurrence reflects a restart penalty similar to a branch mispredict. This is relatively rare.

PMCx0C8 Retired Near Returns

PERF_CTL[5:0]. The number of near return instructions (RET or RET Iw) retired.

PMCx0C9 Retired Near Returns Mispredicted

PERF_CTL[5:0]. The number of near returns retired that were not correctly predicted by the return address predictor. Each such mispredict incurs the same penalty as a mispredicted conditional branch instruction.

PMCx0CA Retired Indirect Branches Mispredicted

PERF_CTL[5:0]. The number of indirect branch instructions retired where the target address was not correctly predicted.

PMCx0CB Retired MMX/FP Instructions

PERF_CTL[5:0]. The number of MMX, SSE or x87 instructions retired. The UnitMask allows the selection of the individual classes of instructions as given in the table. Each increment represents one complete instruction. Since this event includes non-numeric instructions it is not suitable for measuring MFLOPS.

UnitMask	Description
7:3	Reserved.
2	SSE instructions (SSE, SSE2, SSE3, SSSE3, SSE4A, SSE4.1, SSE4.2, AVX, XOP, FMA4)
1	MMX™ instructions.
0	x87 instructions.

PMCx0CD Interrupts-Masked Cycles

PERF_CTL[5:0]. The number of processor cycles where interrupts are masked (EFLAGS.IF = 0). Using edge-counting with this event gives the number of times IF is cleared; dividing the cycle-count value by this value gives the average length of time that interrupts are disabled on each instance. Compare the edge count with

PMCx0CF to determine how often interrupts are disabled for interrupt handling vs. other reasons (e.g. critical sections).

PMCx0CE Interrupts-Masked Cycles with Interrupt Pending

PERF_CTL[5:0]. The number of processor cycles where interrupts are masked (EFLAGS.IF = 0) and an interrupt is pending. Using edge-counting with this event and comparing the resulting count with the edge count for **PMCx0CD** gives the proportion of interrupts for which handling is delayed due to prior interrupts being serviced, critical sections, etc. The cycle count value gives the total amount of time for such delays. The cycle count divided by the edge count gives the average length of each such delay.

PMCx0CF Interrupts Taken

PERF_CTL[5:0]. The number of hardware interrupts taken. This does not include software interrupts (INT n instruction).

PMCx0D0 Decoder Empty

PERF_CTL[2:0]. The number of processor cycles where the decoder has nothing to dispatch (typically waiting on an instruction fetch that missed the Icache, or for the target fetch after a branch mispredict).

PMCx0D1 Dispatch Stalls

PERF_CTL[2:0]. The number of processor cycles where the decoder is stalled for any reason (has one or more instructions ready but can't dispatch them due to resource limitations in execution). This event requires that the other core of the compute unit is in the Halt state. This is the combined effect of events D2h - DAh, some of which may overlap; this event reflects the net stall cycles. The more common stall conditions (events **PMCx0D5**, **PMCx0D6**, **PMCx0D7**, **PMCx0D8**) may overlap considerably. The occurrence of these stalls is highly dependent on the nature of the code being executed (instruction mix, memory reference patterns, etc.).

PMCx0D3 Microsequencer Stall due to Serialization

PERF_CTL[2:0]. The number of processor cycles the micro-sequencer is stalled due to a serializing operation, which waits for the execution pipeline to drain. Relatively rare; mainly associated with system instructions. See **PMCx0D1**.

PMCx0D5 Dispatch Stall for Instruction Retire Q Full

PERF_CTL[2:0]. The number of processor cycles the decoder is stalled because the instruction retire Q is full. This event requires that the other core of the compute unit is in the Halt state. May occur simultaneously with certain other stall conditions; see **PMCx0D1**.

PMCx0D6 Dispatch Stall for Integer Scheduler Queue Full

PERF_CTL[2:0]. The number of processor cycles the decoder is stalled because a required integer unit scheduler queue is full. This event requires that the other core of the compute unit is in the Halt state. May occur simultaneously with certain other stall conditions; see **PMCx0D1**.

PMCx0D7 Dispatch Stall for FP Scheduler Queue Full

PERF_CTL[2:0]. The number of processor cycles the decoder is stalled because the scheduler for the Floating Point scheduler queue is full. This event requires that the other core of the compute unit is in the Halt state. This condition can be caused by a lack of parallelism in FP-intensive code, or by cache misses on FP operand loads (which could also show up as **PMCx0D8** instead, depending on the nature of the instruction sequences). May occur simultaneously with certain other stall conditions; see **PMCx0D1**.

PMCx0D8 Dispatch Stall for LDQ Full

PERF_CTL[2:0]. The number of processor cycles the decoder is stalled because the load queue is full. This event requires that the other core of the compute unit is in the Halt state. This generally occurs due to heavy cache miss activity. May occur simultaneously with certain other stall conditions; see **PMCx0D1**.

PMCx0D9 Microsequencer Stall Waiting for All Quiet

PERF_CTL[2:0]. The number of processor cycles the micro-sequencer is stalled waiting for all outstanding requests to the system to be resolved. Relatively rare; associated with certain system instructions and types of interrupts. May partially overlap certain other stall conditions; see **PMCx0D1**.

PMCx0DB FPU Exceptions

PERF_CTL[5:0]. The number of floating point unit exceptions for microcode assists. The UnitMask may be used to isolate specific types of exceptions.

UnitMask	Description
7:5	Reserved.
4	Bypass faults
3	Ext2Int faults
2	Int2Ext faults
1	Total microtraps
0	Total microfaults

PMCx0D[F:C] DR[3:0] Breakpoint Matches

Table 246: Register Mapping for **PMCx0D[F:C]**

Register	Function
PMCx0DC	DR0
PMCx0DD	DR1
PMCx0DE	DR2
PMCx0DF	DR3

PERF_CTL[5:0]. The number of matches on the address in breakpoint register DR[3:0], per the breakpoint type specified in DR7. Matches occur if the access becomes non-speculative, but not necessarily retired. Each instruction breakpoint match incurs an overhead of about 120 cycles; load/store breakpoint matches do not incur any overhead.

PMCx1C0 Retired x87 Floating Point Operations

PERF_CTL[5:3]. The number of x87 floating point ops that have retired.

UnitMask	Description
7:3	Reserved.
2	Divide and square root ops
1	Multiply ops
0	Add/subtract ops

PMCx1CF Tagged IBS Ops

PERF_CTL[5:0].

UnitMask	Description
7:3	Reserved.
2	Number of times an op could not be tagged by IBS because of a previous tagged op that has not retired.
1	Number of ops tagged by IBS that retired.
0	Number of ops tagged by IBS.

PMCx1D8 Dispatch Stall for STQ Full

PERF_CTL[5:0]. The number of processor cycles the decoder is stalled because the store queue is full. This generally occurs due to heavy cache miss activity. This event requires that the other core of the compute unit is in the Halt state. May occur simultaneously with certain other stall conditions.

3.23 NB Performance Counter Events

This section provides the performance counter events that may be selected through [MSRC001_024\[6,4,2,0\]](#) [Northbridge Performance Event Select (NB_PERF_CTL[3:0])][EventSelect and UnitMask]. See that register and [MSRC001_024\[7,5,3,1\]](#) [Northbridge Performance Event Counter (NB_PERF_CTR[3:0])].

3.23.1 0E[7:0] Events (Memory Controller)

NBPMCx0E0 DRAM Accesses

The number of memory accesses performed by the local DRAM controller. UnitMask[7:0] may be used to isolate the different DRAM page access cases. Page miss cases incur an extra latency to open a page; page conflict cases incur both a page-close as well as page-open penalties. These penalties may be overlapped by DRAM accesses for other requests and don't necessarily represent lost DRAM bandwidth. The associated penalties are as follows:

Page miss: Trcd (DRAM RAS-to-CAS delay)
 Page conflict: Trp + Trcd (DRAM row-precharge time plus RAS-to-CAS delay)

Each DRAM access represents one 64-byte block of data transferred if the DRAM is configured for 64-byte granularity, or one 32-byte block if the DRAM is configured for 32-byte granularity. (The latter is only applicable to single-channel DRAM systems, which may be configured either way.)

UnitMask	Description
7:6	Reserved.
5	DCT1 Page Conflict
4	DCT1 Page Miss
3	DCT1 Page hit
2	DCT0 Page Conflict
1	DCT0 Page Miss
0	DCT0 Page hit

NBPMCx0E1 DRAM Controller Page Table Overflows

The number of page table overflows in the local DRAM controller. This table maintains information about which DRAM pages are open. An overflow occurs when a request for a new page arrives when the maximum number of pages are already open. Each occurrence reflects an access latency penalty equivalent to a page conflict.

UnitMask	Description
7:2	Reserved.
1	DCT1 Page Table Overflow
0	DCT0 Page Table Overflow

NBPMCx0E2 Memory Controller DRAM Command Slots Missed

UnitMask	Description
7:2	Reserved.
1	DCT1 Command Slots Missed (in MEMCLKs)
0	DCT0 Command Slots Missed (in MEMCLKs)

NBPMCx0E3 Memory Controller Turnarounds

The number of turnarounds on the local DRAM data bus. UnitMask[7:0] may be used to isolate the different cases. These represent lost DRAM bandwidth, which may be calculated as follows (in bytes per occurrence):

DIMM turnaround: $\text{DRAM_width_in_bytes} * 2 \text{ edges_per_memclk} * 2$

R/W turnaround: $\text{DRAM_width_in_bytes} * 2 \text{ edges_per_memclk} * 1$

R/W turnaround: $\text{DRAM_width_in_bytes} * 2 \text{ edges_per_memclk} * (\text{Tcl}-1)$

where DRAM_width_in_bytes is 8 or 16 (for single- or dual-channel systems), and Tcl is the CAS latency of the DRAM in memory system clock cycles (where the memory clock for DDR-400, or PC3200 DIMMS, for example, would be 200 MHz).

UnitMask	Description
7:6	Reserved.
5	DCT1 Write to read turnaround
4	DCT1 Read to write turnaround
3	DCT1 DIMM (chip select) turnaround
2	DCT0 Write to read turnaround
1	DCT0 Read to write turnaround
0	DCT0 DIMM (chip select) turnaround

NBPMCx0E4 Memory Controller Bypass Counter Saturation

UnitMask	Description
7:4	Reserved.
3	DCT1 DCQ bypass
2	DCT0 DCQ bypass
1	Memory controller medium priority bypass
0	Memory controller high priority bypass

3.23.2 0E[F:8] Events (Crossbar)

NBPMCx0E8 Thermal Status

UnitMask	Description
7	Reserved
6	Number of clocks HTC P-state is active
5	Number of clocks HTC P-state is inactive
4	Reserved.
3	Reserved.
2	Number of times the HTC trip point is crossed
1:0	Reserved

NBPMCx0E9 CPU/IO Requests to Memory/IO

These events reflect request flow between units and nodes, as selected by the UnitMask. The UnitMask is divided into two fields: request type (CPU or IO access to IO or Memory) and source/target location (local vs. remote). One or more requests types must be enabled via bits 3:0, and at least one source and one target location must be selected via bits 7:4. Each event reflects a request of the selected type(s) going from the selected source(s) to the selected target(s).

Not all possible paths are supported. The following table shows the UnitMask values that are valid for each request type:

Source/Target	CPU to Mem	CPU to IO	IO to Mem	IO to IO
Local -> Local	A8h	A4h	A2h	A1h
Local -> Remote	98h	94h	92h	91h
Remote -> Local	-	64h	-	61h
Remote -> Remote	-	-	-	-

Any of the mask values shown may be logically ORed to combine the events. For instance, local CPU requests to both local and remote nodes would be $A8h | 98h = B8h$. Any CPU to any IO would be $A4h | 94h | 64h = F4h$ (but remote CPU to remote IO requests would not be included).

Note: It is not possible to tell from these events how much data is going in which direction, as there is no distinction between reads and writes. Also, particularly for IO, the requests may be for varying amounts of data, anywhere from one to sixty-four bytes. For a direct measure of the amount and direction of data flowing between nodes, use events F6h, F7h and F8h.

UnitMask	Description
7	From local node
6	From remote node
5	To local node
4	To remote node

3	CPU to Mem
2	CPU to IO
1	IO to Mem
0	IO to IO

NBPMCx0EA Cache Block Commands

The number of requests made to the system for cache line transfers or coherency state changes, by request type. Each increment represents one cache line transfer, except for Change-to-Dirty. If a Change-to-Dirty request hits on a line in another processor's cache that's in the Owned state, it causes a cache line transfer, otherwise there is no data transfer associated with Change-to-Dirty requests.

UnitMask	Description
7:6	Reserved.
5	Change-to-Dirty (first store to clean block already in cache)
4	Read Block Modified (Dcache store miss refill)
3	Read Block Shared (Icache refill)
2	Read Block (Dcache load miss refill)
1	Reserved.
0	Victim Block (Writeback)

NBPMCx0EB Sized Commands

The number of Sized Read/Write commands handled by the System Request Interface (local processor and hostbridge interface to the system). These commands may originate from the processor or hostbridge. Typical uses of the various Sized Read/Write commands are given in the UnitMask table. See [NBPMCx0EC](#), which provides a separate measure of Hostbridge accesses.

UnitMask	Description
7:6	Reserved.
5	SzRd DW (1-16 DWORDs). Typical Usage: Block-oriented DMA reads, typically cache-line size.
4	SzRd Byte (4 bytes). Typical Usage: Legacy or mapped IO.
3	Posted SzWr DW (1-16 DWORDs). Typical Usage: Block-oriented DMA writes, often cache-line sized; also processor Write Combining buffer flushes.
2	Posted SzWr Byte (1-32 bytes). Typical Usage: Sub-cache-line DMA writes, size varies; also flushes of partially-filled Write Combining buffer.
1	Non-Posted SzWr DW (1-16 DWORDs). Typical Usage: Legacy or mapped IO, typically 1 DWORD.
0	Non-Posted SzWr Byte (1-32 bytes). Typical Usage: Legacy or mapped IO, typically 1-4 bytes.

NBPMCx0EC Probe Responses and Upstream Requests

This covers two unrelated sets of events: cache probe results, and requests received by the hostbridge from devices on non-coherent links.

Probe results: These events reflect the results of probes sent from a memory controller to local caches. They provide an indication of the degree data and code is shared between processors (or moved between processors due to process migration). The dirty-hit events indicate the transfer of a 64-byte cache line to the requestor (for a read or cache refill) or the target memory (for a write). The system bandwidth used by these, in terms of bytes per unit of time, may be calculated as 64 times the event count, divided by the elapsed time. Sized writes to memory that cover a full cache line do not incur this cache line transfer -- they simply invalidate the line and are reported as clean hits. Cache line transfers occur for Change2Dirty requests that hit cache lines in the Owned state. (Such cache lines are counted as Modified-state refills for [PMCx06C](#), System Read Responses.)

Upstream requests: The upstream read and write events reflect requests originating from a device on a local IO link. The two read events allow display refresh traffic in a UMA system to be measured separately from other DMA activity. Display refresh traffic is typically dominated by 64-byte transfers. Non-display-related DMA accesses may be anywhere from 1 to 64 bytes in size, but may be dominated by a particular size such as 32 or 64 bytes, depending on the nature of the devices.

UnitMask	Description
7	Upstream non-ISOC writes
6	Upstream ISOC writes
5	Upstream non-display refresh reads
4	Upstream display refresh/ISOC reads
3	Probe hit dirty with memory cancel (probed by DMA read or cache refill request)
2	Probe hit dirty without memory cancel (probed by Sized Write or Change2Dirty)
1	Probe hit clean
0	Probe miss

3.23.3 0F[F:0] Events (Crossbar)

3.23.4 1E[F:0] Events (Crossbar)

NBPMCx1E0 CPU to DRAM Requests to Target Node

This event counts all DRAM reads and writes generated by cores on the local node to the targeted node in the coherent fabric. This counter can be used to observe processor data affinity in NUMA aware operating systems.

UnitMask	Description
7	From Local node to Node 7
6	From Local node to Node 6
5	From Local node to Node 5
4	From Local node to Node 4
3	From Local node to Node 3
2	From Local node to Node 2
1	From Local node to Node 1
0	From Local node to Node 0

NBPMCx1E1 IO to DRAM Requests to Target Node

This event counts all DRAM reads and writes generated by IO devices attached to the IO links of the local node the targeted node in the coherent fabric. This counter can be used to observe IO device data affinity in NUMA aware operating systems.

UnitMask	Description
7	From Local node to Node 7
6	From Local node to Node 6
5	From Local node to Node 5
4	From Local node to Node 4
3	From Local node to Node 3
2	From Local node to Node 2
1	From Local node to Node 1
0	From Local node to Node 0

NBPMCx1E2 CPU Read Command Latency to Target Node 0-3

This event counts the number of NB clocks from when the targeted command is received in the NB to when the targeted command completes. This event only tracks one outstanding command at a time. To determine latency between the local node and a remote node set UnitMask[7:4] to select the node and UnitMask[3:0] to select the read command type. The count returned by the counter should be divided by the count returned by [NBPMCx1E3](#) do determine the average latency for the command type.

UnitMask	Description
7	From Local node to Node 3
6	From Local node to Node 2
5	From Local node to Node 1
4	From Local node to Node 0
3	Change-to-Dirty
2	Read block modified
1	Read block shared
0	Read block

NBPMCx1E3 CPU Read Command Requests to Target Node 0-3

This event counts the number of requests that a latency measurement is made for using [NBPMCx1E2](#). To determine the number of commands that a latency measurement are made for between the local node and a remote node set UnitMask[7:4] to select the node and UnitMask[3:0] to select the read command type.

UnitMask	Description
7	From Local node to Node 3
6	From Local node to Node 2
5	From Local node to Node 1
4	From Local node to Node 0

3	Change-to-Dirty
2	Read block modified
1	Read block shared
0	Read block

NBPMCx1E4 CPU Read Command Latency to Target Node 4-7

This event counts the number of NB clocks from when the targeted command is received in the NB to when the targeted command completes. This event only tracks one outstanding command at a time. To determine latency between the local node and a remote node set UnitMask[7:4] to select the node and UnitMask[3:0] to select the read command type. The count returned by the counter should be divided by the count returned by [NBPMCx1E5](#) do determine the average latency for the command type.

UnitMask	Description
7	From Local node to Node 7
6	From Local node to Node 6
5	From Local node to Node 5
4	From Local node to Node 4
3	Change-to-Dirty
2	Read block modified
1	Read block shared
0	Read block

NBPMCx1E5 CPU Read Command Requests to Target Node 4-7

This event counts the number of requests that a latency measurement is made for using [NBPMCx1E4](#). To determine the number of commands that a latency measurement are made for between the local node and a remote node set UnitMask[7:4] to select the node and UnitMask[3:0] to select the read command type.

UnitMask	Description
7	From Local node to Node 7
6	From Local node to Node 6
5	From Local node to Node 5
4	From Local node to Node 4
3	Change-to-Dirty
2	Read block modified
1	Read block shared
0	Read block

NBPMCx1E6 CPU Command Latency to Target Node 0-3/4-7

This event counts the number of NB clocks from when the targeted command is received in the NB to when the targeted command completes. This event only tracks one outstanding command at a time. To determine latency between the local node and a remote node set UnitMask[7:4] to select the node, UnitMask[3] to select the node group and UnitMask[3:0] to select the command type. The count returned by the counter should be divided by

the count returned by [NBPMCx1E7](#) do determine the average latency for the command type.

UnitMask	Description
7	From Local node to Node 3/7
6	From Local node to Node 2/6
5	From Local node to Node 1/5
4	From Local node to Node 0/4
3	Node Group Select. 0=Nodes 0-3. 1= Nodes 4-7.
2	Victim Block
1	Write Sized
0	Read Sized

NBPMCx1E7 CPU Requests to Target Node 0-3/4-7

This event counts the number of requests that a latency measurement is made for using [NBPMCx1E6](#). To determine the number of commands that a latency measurement are made for between the local node and a remote node set UnitMask[7:4] to select the node, UnitMask[3] to select the node group and UnitMask[3:0] to select the command type.

UnitMask	Description
7	From Local node to Node 3/7
6	From Local node to Node 2/6
5	From Local node to Node 1/5
4	From Local node to Node 0/4
3	Node Group Select. 0=Nodes 0-3. 1= Nodes 4-7.
2	Victim Block
1	Write Sized
0	Read Sized

NBPMCx1EB Request Cache Status 1

The probe response type for RdBlkM or ChgToDirty request type.

UnitMask	Description
7	Track Cache Stat for RdBlkM
6	Track Cache Stat for ChgToDirty
5	Directed Probe
4	Probe Miss
3	Probe Hit M
2	Probe Hit MuW or O
1	Probe Hit E
0	Probe Hit S

3.23.5 1F[F:0] Events (Memory Controller, Crossbar)

NBPMCx1F0 Memory Controller Requests

Read/Write requests: The read/write request events reflect the total number of commands sent to the DRAM controller.

Sized Read/Write activity: The Sized Read/Write events reflect 32- or 64-byte transfers (as opposed to other sizes which could be anywhere between 1 and 64 bytes), from either the processor or the Hostbridge (on any node in an MP system). Such accesses from the processor would be due only to write combining buffer flushes, where 32-byte accesses would reflect flushes of partially-filled buffers. [PMCx065](#) provides a count of sized write requests associated with WC buffer flushes; comparing that with counts for these events (providing there is very little Hostbridge activity at the same time) gives an indication of how efficiently the write combining buffers are being used. [PMCx065](#) may also be useful in factoring out WC flushes when comparing these events with the Upstream Requests component of [PMCx06C](#).

UnitMask	Description
7	Read requests sent to the DCT while writes requests are pending in the DCT
6	64 Byte Sized Reads
5	32 Bytes Sized Reads
4	64 Bytes Sized Writes
3	32 Bytes Sized Writes
2	Prefetch requests sent to the DCT
1	Read requests (including prefetch requests) sent to the DCT
0	Write requests sent to the DCT

4 Register List

The following is a list of all storage elements, context, and registers provided in this document. Page numbers, register mnemonics, and register names are provided.

42	SMMFEC0: SMM IO Trap Offset	182	D0F0xBC_x1F2[E8:08:step20]: LCLK DPM Control 2
43	SMMFEC4: Local SMI Status	182	D0F0xBC_x1F2[F0:10:step20]: LCLK DPM Activity Thresholds
43	SMMFEC8: SMM IO Restart Byte	183	D0F0xBC_x1F300: SMU_LCLK_DPM_CNTL
44	SMMFEC9: Auto Halt Restart Offset	183	D0F0xBC_x1F308:
44	SMMFECA: NMI Mask		SMU_LCLK_DPM_THERMAL_THROTTLING_CNTL
44	SMMFED8: SMM SVM State	183	D0F0xBC_x1F30C:
45	SMMFEFC: SMM-Revision Identifier		SMU_LCLK_DPM_THERMAL_THROTTLING_THRESHOLDS
45	SMMFF00: SMM Base Address (SMM_BASE)	184	D0F0xBC_x1F380: FIRMWARE_FLAGS
167	IOCF8: IO-Space Configuration Address	184	D0F0xBC_x1F384: FIRMWARE_VID
167	IOCF8: IO-Space Configuration Data Port	184	D0F0xBC_x1F388: TEMPERATURE_READ_ADDR
168	D0F0x00: Device/Vendor ID	184	D0F0xBC_x1F39C: PCIE_PG_ARGS
168	D0F0x04: Status/Command	185	D0F0xBC_x1F3D8: LOAD_LINE_TRIM_TABLE1
169	D0F0x08: Class Code/Revision ID	185	D0F0xBC_x1F3DC: LOAD_LINE_TRIM_TABLE2
169	D0F0x0C: Header Type	185	D0F0xBC_x1F3F8: CSR_GNB_1
169	D0F0x2C: Subsystem and Subvendor ID	186	D0F0xBC_x1F3FC: CSR_GNB_2
169	D0F0x34: Capabilities Pointer	186	D0F0xBC_x1F400: CSR_GNB_3
169	D0F0x4C: PCI Control	186	D0F0xBC_x1F404: LOAD_LINE_OFFSET_TABLE
170	D0F0x60: Miscellaneous Index	186	D0F0xBC_x1F428: PM_CONFIG
170	D0F0x64: Miscellaneous Index Data	188	D0F0xBC_x1F460: PM_INTERVAL_CNTL_0
170	D0F0x64_x00: Northbridge Control	188	D0F0xBC_x1F468: PM_TIMER_PERIOD
171	D0F0x64_x0B: IOC Link Control	188	D0F0xBC_x1F46C: PM_TIMERS_1
171	D0F0x64_x0C: IOC Bridge Control	188	D0F0xBC_x1F5F8: NB_PSTATE_CONFIG
171	D0F0x64_x0D: IOC PCI Configuration	189	D0F0xBC_x1F5FC: NB_PSTATE_STATUS
172	D0F0x64_x16: IOC Advanced Error Reporting Control	189	D0F0xBC_x1F628: LHTC CONFIG 1
172	D0F0x64_x19: Top of Memory 2 Low	189	D0F0xBC_x1F62C: TDC_VRM_LIMIT
172	D0F0x64_x1A: Top of Memory 2 High	189	D0F0xBC_x1F638: PM_TIMERS_2
172	D0F0x64_x1C: Internal Graphics PCI Control 1	190	D0F0xBC_x1F86C: LHTC CONFIG 2
173	D0F0x64_x1D: Internal Graphics PCI Control 2	190	D0F0xBC_x1F89C: LHTC CONFIG 3
174	D0F0x64_x20: Programmable Device Remap 0	190	D0F0xBC_x1F8EC: HYBRID_BOOST_CONFIG
174	D0F0x64_x21: Programmable Device Remap 1	191	D0F0xBC_xE000120: Activity Monitor Control
174	D0F0x64_x22: LCLK Control 0	191	D0F0xBC_xE0003000: CPU Interrupt Request
175	D0F0x64_x23: LCLK Control 1	191	D0F0xBC_xE0003004: CPU Interrupt Status
175	D0F0x64_x46: IOC Features Control	192	D0F0xBC_xE0003048: SCLK_MIN_DIV
176	D0F0x64_x5[B,9,7,5,3,1]: IOC PCIe Device Control	192	D0F0xBC_xE0003088: SMU_AUTH_STATUS
176	D0F0x7C: IOC Configuration Control	192	D0F0xBC_xE00030A4: SMU_FIRMWARE_AUTH
176	D0F0x84: Link Arbitration	192	D0F0xBC_xE0104168: MEMCLK VID Configuration
177	D0F0x90: Northbridge Top of Memory	193	D0F0xBC_xE010416C: MEMCLK VID Configuration
177	D0F0x94: Northbridge ORB Configuration Offset	193	D0F0xBC_xE0104170: MEMCLK VID Configuration
177	D0F0x98: Northbridge ORB Configuration Data Port	193	D0F0xBC_xE0104184: SVI Loadline Configuration
177	D0F0x98_x06: ORB Downstream Control 0	193	D0F0xBC_xE0104188: SMU Power Management Config 1
177	D0F0x98_x07: ORB Upstream Arbitration Control 0	194	D0F0xBC_xE010418C: SMU Power Management Config 2
178	D0F0x98_x08: ORB Upstream Arbitration Control 1	194	D0F0xBC_xE010703C: SMU Power Management Config 3
178	D0F0x98_x09: ORB Upstream Arbitration Control 2	194	D0F0xBC_xE0300000: SBLK_MC_PGFSM_CONFIG
178	D0F0x98_x0C: ORB Upstream Arbitration Control 5	194	D0F0xBC_xE030000C: SBLK_PGD_PGFSM_CONFIG
179	D0F0x98_x1E: ORB Receive Control 0	195	D0F0xBC_xE0300018: SBLK_IOMMU_PGFSM_CONFIG
179	D0F0x98_x26: ORB IOMMU Control 0	195	D0F0xBC_xE0300024: SBLK_VCE_PGFSM_CONFIG
180	D0F0x98_x27: ORB IOMMU Control 1	195	D0F0xBC_xE0300030: SBLK_DC2_PGFSM_CONFIG
180	D0F0x98_x28: ORB Transmit Control 0	196	D0F0xBC_xE030003C: UVD_CHAIN_PGFSM_CONFIG
180	D0F0x98_x2C: ORB Clock Control	196	D0F0xBC_xE0300200: SBLK_MC_PGFSM_DEBUG
180	D0F0x98_x3A: ORB Source Tag Translation Control 2	196	D0F0xBC_xE0300208: SBLK_IOMMU_PGFSM_DEBUG
181	D0F0x98_x4[A,9]: ORB LCLK Clock Control 1-0	196	D0F0xBC_xE030020C: SBLK_VCE_PGFSM_DEBUG
181	D0F0xB8: SMU Index Address	197	D0F0xBC_xE0300210: SBLK_DC2_PGFSM_DEBUG
181	D0F0xBC: SMU Index Data	197	D0F0xBC_xE0300218: SBLK_UVDU_PGFSM_DEBUG
181	D0F0xBC_x1F100: SCLK DPM Control	197	D0F0xBC_xE0300320: PGFSM_CLK_CNTL_0
182	D0F0xBC_x1F2[E0:00:step20]: LCLK DPM Control 0	197	D0F0xBC_xE0300324: PGFSM_CLK_CNTL_1

197 D0F0xE0: Link Index Address
 198 D0F0xE4: Link Index Data
 198 D0F0xE4_x0[2:1]01_0002: IO Link Hardware Debug
 198 D0F0xE4_x0[2:1]01_0010: IO Link Control 1
 199 D0F0xE4_x0[2:1]01_0011: IO Link Config Control
 200 D0F0xE4_x0[2:1]01_001C: IO Link Control 2
 200 D0F0xE4_x0[2:1]01_0020: IO Link Chip Interface Control
 200 D0F0xE4_x0[2:1]01_0040: IO Link Phy Control
 201 D0F0xE4_x0[2:1]01_00B0: IO Link Strap Link Strap Control
 201 D0F0xE4_x0[2:1]01_00C0: IO Link Strap Miscellaneous
 201 D0F0xE4_x0[2:1]01_00C1: IO Link Strap Miscellaneous
 201 D0F0xE4_x0[2:1]1[3:0]_0010: PIF Control
 202 D0F0xE4_x0[2:1]1[3:0]_0011: PIF Pairing
 203 D0F0xE4_x0[2:1]1[3:0]_001[3:2]: PIF Power Down Control [1:0]
 204 D0F0xE4_x0[2:1]1[3:0]_0015: PIF Transmitter Status
 205 D0F0xE4_x0[2:1]2[3:0]_0000: Phy Compensation Control and Calibration Control I
 205 D0F0xE4_x0[2:1]2[3:0]_000[2:1]: Phy Impedance Control
 207 D0F0xE4_x0[2:1]2[3:0]_000[A:9]: Phy Clock Tree Control
 208 D0F0xE4_x0[2:1]2[3:0]_000[C:B]: Phy Serial Bus Packet Control
 209 D0F0xE4_x0[2:1]2[3:0]_000D: Phy Serial Bus Compensation Component Packet Enable
 209 D0F0xE4_x0[2:1]2[3:0]_2000: Phy PLL Power State Control
 210 D0F0xE4_x0[2:1]2[3:0]_2002: Phy PLL Control
 210 D0F0xE4_x0[2:1]2[3:0]_2005: Phy PLL Frequency and Mode Control
 211 D0F0xE4_x0[2:1]2[3:0]_2008: Phy PLL Update Control
 213 D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]1: Phy Receiver DLL Control and Test 1
 213 D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]2: Phy Receiver Phase Loop Filter Control
 215 D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]5: Phy Receiver Timing Margin Test
 216 D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]6: Phy Receiver DFE and DFR Control
 217 D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]A: Phy DLL Test and Control 3
 219 D0F0xE4_x0[2:1]2[1:0]_[5:4][7:6,3:0][8,0]F: Phy Receiver DLL Test and Debug 5
 220 D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]0: Phy Tx Deemphasis and Margining Control
 221 D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]5: Phy Transmit Link Configuration
 222 D0F0xE4_x0[2:1]2[3:0]_[7:6][7:6,3:0][8,0]6: Phy Transmit Nominal Deemphasis Control
 223 D0F0xE4_x0[2:1]2[1:0]_[D:C][7:0][8,0]5: Termination Mode Control
 224 D0F0xE4_x013[1:0]_0046: Subsystem and Vendor ID
 224 D0F0xE4_x013[1:0]_0080: Link Configuration
 225 D0F0xE4_x013[1:0]_0[C:8]00: Link Training Control
 225 D0F0xE4_x013[1:0]_0[C:8]03: Link Deemphasis Control
 225 D0F0xE4_x013[3:0]_8011: Link Transmit Clock Gating Control
 226 D0F0xE4_x013[3:0]_8012: Link Idle-Resume Clock Gating Control
 226 D0F0xE4_x013[3:0]_8013: Transmit Clock Pll Control
 227 D0F0xE4_x013[3:0]_8014: Link Transmit Clock Gating Control 2
 229 D0F0xE4_x013[3:0]_8015: IO Link IOC Control
 229 D0F0xE4_x013[3:0]_8016: Link Clock Switching Control
 230 D0F0xE4_x013[3:0]_8020: Lane Control
 230 D0F0xE4_x013[3:0]_8021: Transmitter Lane Mux
 230 D0F0xE4_x013[3:0]_8022: Receiver Lane Mux
 231 D0F0xE4_x013[3:0]_8023: Lane Enable
 231 D0F0xE4_x013[3:0]_8025: Lane Mux Power Sequence Control
 232 D0F0xE4_x013[1:0]_8031: Lane Counter Status
 232 D0F0xE4_x013[3:0]_8060: Soft Reset Command 0
 232 D0F0xE4_x013[3:0]_8062: Soft Reset Control 0
 233 D0F0xE4_x013[3:0]_8063: Soft Reset Control 1
 233 D0F0xE4_x0130_80F0: BIOS Timer
 234 D0F0xE4_x0130_80F1: BIOS Timer Control
 234 D0F0xE4_x013[3:1]_804[3:0]: DDI Slice
 234 D0F0xE4_x013[3:1]_804[E:8]: DDI Dig
 236 D0F2x00: Device/Vendor ID
 236 D0F2x04: Status/Command
 237 D0F2x08: Class Code/Revision ID
 237 D0F2x0C: Header Type
 237 D0F2x2C: Subsystem and Subvendor ID
 237 D0F2x34: Capabilities Pointer

237 D0F2x3C: Interrupt Line
 238 D0F2x40: IOMMU Capability
 238 D0F2x44: IOMMU Base Address Low
 239 D0F2x48: IOMMU Base Address High
 239 D0F2x4C: IOMMU Range
 239 D0F2x50: IOMMU Miscellaneous Information Register
 240 D0F2x54: IOMMU MSI Capability Register
 240 D0F2x58: IOMMU MSI Address Low
 241 D0F2x5C: IOMMU MSI Address High
 241 D0F2x60: IOMMU MSI Data
 241 D0F2x64: IOMMU MSI Mapping Capability
 241 D0F2x6C: IOMMU Control
 241 D0F2x70: IOMMU MMIO Control Low
 242 D0F2x74: IOMMU MMIO Control High
 242 D0F2x78: IOMMU Range Control
 242 D0F2xF0: IOMMU L2 Config Index
 243 D0F2xF4: IOMMU L2 Config Data
 243 D0F2xF4_x00: L2_PERF_CNTL_0
 243 D0F2xF4_x01: L2_PERF_COUNT_0
 243 D0F2xF4_x02: L2_PERF_COUNT_1
 243 D0F2xF4_x03: L2_PERF_CNTL_1
 243 D0F2xF4_x04: L2_PERF_COUNT_2
 244 D0F2xF4_x05: L2_PERF_COUNT_3
 244 D0F2xF4_x06: L2_DEBUG_0
 244 D0F2xF4_x07: L2_DEBUG_1
 244 D0F2xF4_x08: L2_STATUS_0
 244 D0F2xF4_x0C: L2_CONTROL_0
 245 D0F2xF4_x0D: L2_CONTROL_1
 245 D0F2xF4_x10: L2_DTC_CONTROL
 246 D0F2xF4_x11: L2_DTC_HASH_CONTROL
 246 D0F2xF4_x12: L2_DTC_WAY_CONTROL
 247 D0F2xF4_x14: L2_ITC_CONTROL
 247 D0F2xF4_x15: L2_ITC_HASH_CONTROL
 248 D0F2xF4_x16: L2_ITC_WAY_CONTROL
 248 D0F2xF4_x18: L2_PTC_A_CONTROL
 249 D0F2xF4_x19: L2_PTC_A_HASH_CONTROL
 249 D0F2xF4_x1A: L2_PTC_A_WAY_CONTROL
 249 D0F2xF4_x20: L2_CREDIT_CONTROL_2
 250 D0F2xF4_x22: L2A_UPDATE_FILTER_CNTL
 250 D0F2xF4_x30: L2_ERR_RULE_CONTROL_3
 250 D0F2xF4_x31: L2_ERR_RULE_CONTROL_4
 250 D0F2xF4_x32: L2_ERR_RULE_CONTROL_5
 250 D0F2xF4_x33: L2_L2A_CK_GATE_CONTROL
 251 D0F2xF4_x34: L2_L2A_PGSIIZE_CONTROL
 251 D0F2xF4_x40: L2_PERF_CNTL_2
 251 D0F2xF4_x41: L2_PERF_COUNT_4
 251 D0F2xF4_x42: L2_PERF_COUNT_5
 252 D0F2xF4_x43: L2_PERF_CNTL_3
 252 D0F2xF4_x44: L2_PERF_COUNT_6
 252 D0F2xF4_x45: L2_PERF_COUNT_7
 252 D0F2xF4_x46: L2_DEBUG_2
 252 D0F2xF4_x47: L2_DEBUG_3
 252 D0F2xF4_x48: L2_STATUS_1
 253 D0F2xF4_x4C: L2_CONTROL_5
 253 D0F2xF4_x4D: L2_CONTROL_6
 254 D0F2xF4_x50: L2_PDC_CONTROL
 254 D0F2xF4_x51: L2_PDC_HASH_CONTROL
 255 D0F2xF4_x52: L2_PDC_WAY_CONTROL
 255 D0F2xF4_x53: L2B_UPDATE_FILTER_CNTL
 255 D0F2xF4_x54: L2_TW_CONTROL
 256 D0F2xF4_x56: L2_CP_CONTROL
 256 D0F2xF4_x57: L2_CP_CONTROL_1
 256 D0F2xF4_x58: IOMMU_L2_GUEST_ADDR_CNTRL

256 D0F2xF4_x60: L2_TW_CONTROL_1
 257 D0F2xF4_x61: L2_TW_CONTROL_2
 257 D0F2xF4_x62: L2_TW_CONTROL_3
 257 D0F2xF4_x6A: L2_INT_CONTROL
 257 D0F2xF4_x70: L2_CREDIT_CONTROL_0
 258 D0F2xF4_x71: L2_CREDIT_CONTROL_1
 258 D0F2xF4_x78: L2_MCIF_CONTROL
 259 D0F2xF4_x80: L2_ERR_RULE_CONTROL_0
 259 D0F2xF4_x81: L2_ERR_RULE_CONTROL_1
 259 D0F2xF4_x82: L2_ERR_RULE_CONTROL_2
 259 D0F2xF4_x90: L2_L2B_CK_GATE_CONTROL
 260 D0F2xF4_x92: PPR_CONTROL
 260 D0F2xF4_x94: L2_L2B_PGFSIZE_CONTROL
 260 D0F2xF8: IOMMU L1 Config Index
 260 D0F2xFC: IOMMU L1 Config Data
 261 D0F2xFC_x00_L1sel[3:0]: L1_PERF_CNTL
 261 D0F2xFC_x01_L1sel[3:0]: L1_PERF_COUNT_0
 261 D0F2xFC_x02_L1sel[3:0]: L1_PERF_COUNT_1
 261 D0F2xFC_x07_L1sel[3:0]: L1_DEBUG_1
 261 D0F2xFC_x08_L1sel[3:0]: L1_DEBUG_STATUS
 262 D0F2xFC_x0C_L1sel[3:0]: L1_CNTRL_0
 262 D0F2xFC_x0D_L1sel[3:0]: L1_CNTRL_1
 263 D0F2xFC_x0E_L1sel[3:0]: L1_CNTRL_2
 264 D0F2xFC_x0F_L1sel[3:0]: L1_CNTRL_3
 264 D0F2xFC_x10_L1sel[3:0]: L1_BANK_SEL_0
 264 D0F2xFC_x11_L1sel[3:0]: L1_BANK_DISABLE_0
 264 D0F2xFC_x20_L1sel[3:0]: L1_WQ_STATUS_0
 265 D0F2xFC_x21_L1sel[3:0]: L1_WQ_STATUS_1
 265 D0F2xFC_x22_L1sel[3:0]: L1_WQ_STATUS_2
 266 D0F2xFC_x23_L1sel[3:0]: L1_WQ_STATUS_3
 266 D0F2xFC_x32_L1sel[3:0]: L1_CNTRL_4
 266 D0F2xFC_x33_L1sel[3:0]: L1_CLKCNTRL_0
 266 D0F2xFC_x34_L1sel[3:0]: L1_MEMPWRCNTRL_0
 267 D0F2xFC_x35_L1sel[3:0]: L1_MEMPWRCNTRL_1
 267 D0F2xFC_x36_L1sel[3:0]: L1_GUEST_ADDR_CNTRL
 267 D0F2xFC_x37_L1sel[3:0]: L1_FEATURE_SUP_CNTRL
 268 D1F0x00: Device/Vendor ID Register
 268 D1F0x04: Status/Command Register
 269 D1F0x08: Class Code/Revision ID Register
 269 D1F0x0C: Header Type Register
 269 D1F0x10: Graphic Memory Base Address
 269 D1F0x14: Graphics IO Base Address
 270 D1F0x14: Graphics Memory Base Address 64
 270 D1F0x18: Graphics Memory Mapped Registers Base Address
 270 D1F0x1C: Base Address 3
 270 D1F0x1C: Graphics Memory Mapped Registers Address 64
 271 D1F0x20: Base Address 4
 271 D1F0x20: Graphics IO Base Address
 271 D1F0x24: Base Address 5
 271 D1F0x2C: Subsystem and Subvendor ID Register
 271 D1F0x30: Expansion ROM Base Address
 272 D1F0x34: Capabilities Pointer
 272 D1F0x3C: Interrupt Line
 272 D1F0x4C: Subsystem and Subvendor ID Mirror
 272 D1F0x50: Power Management Capability
 273 D1F0x54: Power Management Control and Status
 273 D1F0x58: PCI Express Capability
 273 D1F0x5C: Device Capability
 274 D1F0x60: Device Control and Status
 275 D1F0x64: Link Capability
 275 D1F0x68: Link Control and Status
 276 D1F0x7C: Device Capability 2
 276 D1F0x80: Device Control and Status 2

276	D1F0x88: Link Control and Status 2
277	D1F0xA0: MSI Capability
277	D1F0xA4: MSI Message Address Low
278	D1F0xA8: MSI Message Data
278	D1F0xA8: MSI Message Address High
278	D1F0xAC: MSI Message Data
278	D1F0x100: Vendor Specific Enhanced Capability
278	D1F0x104: Vendor Specific Header
279	D1F0x108: Vendor Specific 1
279	D1F0x10C: Vendor Specific 2
280	D1F1x00: Device/Vendor ID
280	D1F1x04: Status/Command
281	D1F1x08: Class Code/Revision ID
281	D1F1x0C: Header Type
281	D1F1x10: Audio Registers Base Address
281	D1F1x14: Base Address 1
282	D1F1x18: Base Address 2
282	D1F1x1C: Base Address 3
282	D1F1x20: Base Address 4
282	D1F1x24: Base Address 5
282	D1F1x2C: Subsystem and Subvendor ID
282	D1F1x30: Expansion ROM Base Address
283	D1F1x34: Capabilities Pointer
283	D1F1x3C: Interrupt Line
283	D1F1x4C: Subsystem and Subvendor ID Mirror
283	D1F1x50: Power Management Capability
284	D1F1x54: Power Management Control and Status
284	D1F1x58: PCI Express Capability
284	D1F1x5C: Device Capability
285	D1F1x60: Device Control and Status
286	D1F1x64: Link Capability
286	D1F1x68: Link Control and Status
287	D1F1x7C: Device Capability 2
287	D1F1x80: Device Control and Status 2
288	D1F1x84: Link Capability 2
288	D1F1x88: Link Control and Status 2
288	D1F1xA0: MSI Capability
289	D1F1xA4: MSI Message Address Low
289	D1F1xA8: MSI Message Data
289	D1F1xA8: MSI Message Address High
289	D1F1xAC: MSI Message Data
290	D1F1x100: Vendor Specific Enhanced Capability
290	D1F1x104: Vendor Specific Header
290	D1F1x108: Vendor Specific 1
290	D1F1x10C: Vendor Specific 2
291	D[8:2]F0x00: Device/Vendor ID Register
291	D[8:2]F0x04: Status/Command Register
292	D[8:2]F0x08: Class Code/Revision ID Register
292	D[8:2]F0x0C: Header Type Register
292	D[8:2]F0x18: Bus Number and Secondary Latency Register
293	D[8:2]F0x1C: IO Base and Secondary Status Register
293	D[8:2]F0x20: Memory Limit and Base Register
293	D[8:2]F0x24: Prefetchable Memory Limit and Base Register
294	D[8:2]F0x28: Prefetchable Memory Base High Register
294	D[8:2]F0x2C: Prefetchable Memory Limit High Register
294	D[8:2]F0x30: IO Base and Limit High Register
294	D[8:2]F0x34: Capabilities Pointer Register
294	D[8:2]F0x3C: Bridge Control Register
295	D[8:2]F0x50: Power Management Capability Register
295	D[8:2]F0x54: Power Management Control and Status Register
296	D[8:2]F0x58: PCI Express Capability Register
296	D[8:2]F0x5C: Device Capability Register
297	D[8:2]F0x60: Device Control and Status Register

298	D[8:2]F0x64: IO Link Capability Register
299	D[8:2]F0x68: IO Link Control and Status Register
300	D[8:2]F0x6C: Slot Capability Register
301	D[8:2]F0x70: Slot Control and Status Register
302	D[8:2]F0x74: Root Complex Capability and Control Register
302	D[8:2]F0x78: Root Complex Status Register
302	D[8:2]F0x7C: Device Capability 2
303	D[8:2]F0x80: Device Control and Status 2
303	D[8:2]F0x84: IO Link Capability 2
303	D[8:2]F0x88: IO Link Control and Status 2
304	D[8:2]F0x8C: Slot Capability 2
304	D[8:2]F0x90: Slot Control and Status 2
304	D[8:2]F0xA0: MSI Capability Register
305	D[8:2]F0xA4: MSI Message Address Low
305	D[8:2]F0xA8: MSI Message Data
305	D[8:2]F0xA8: MSI Message Address High
305	D[8:2]F0xAC: MSI Message Data
306	D[8:2]F0xB0: Subsystem and Subvendor Capability ID Register
306	D[8:2]F0xB4: Subsystem and Subvendor ID Register
306	D[8:2]F0xB8: MSI Capability Mapping
306	D[8:2]F0xBC: MSI Mapping Address Low
306	D[8:2]F0xC0: MSI Mapping Address High
306	D[8:2]F0xE0: Root Port Index
307	D[8:2]F0xE4: Root Port Data
307	D[8:2]F0xE4_x02: Root Port Hardware Debug
307	D[8:2]F0xE4_x20: Root Port TX Control
307	D[8:2]F0xE4_x50: Root Port Lane Status
308	D[8:2]F0xE4_x6A: Root Port Error Control
308	D[8:2]F0xE4_x70: Root Port Receiver Control
308	D[8:2]F0xE4_xA0: Per Port Link Controller (LC) Control
309	D[8:2]F0xE4_xA1: LC Training Control
309	D[8:2]F0xE4_xA2: LC Link Width Control
310	D[8:2]F0xE4_xA3: LC Number of FTS Control
310	D[8:2]F0xE4_xA4: LC Link Speed Control
311	D[8:2]F0xE4_xA5: LC State 0
312	D[8:2]F0xE4_xB1: LC Control 2
312	D[8:2]F0xE4_xB5: LC Control 3
312	D[8:2]F0xE4_xC0: LC Strap Override
313	D[8:2]F0xE4_xC1: Root Port Miscellaneous Strap Override
313	D[8:2]F0x100: Vendor Specific Enhanced Capability Register
313	D[8:2]F0x104: Vendor Specific Header Register
313	D[8:2]F0x108: Vendor Specific 1 Register
313	D[8:2]F0x10C: Vendor Specific 2 Register
314	D[8:2]F0x128: Virtual Channel 0 Resource Status Register
314	D[8:2]F0x150: Advanced Error Reporting Capability
314	D[8:2]F0x154: Uncorrectable Error Status
315	D[8:2]F0x158: Uncorrectable Error Mask
316	D[8:2]F0x15C: Uncorrectable Error Severity
316	D[8:2]F0x160: Correctable Error Status
317	D[8:2]F0x164: Correctable Error Mask
317	D[8:2]F0x168: Advanced Error Control
317	D[8:2]F0x16C: Header Log DW0
318	D[8:2]F0x170: Header Log DW1
318	D[8:2]F0x174: Header Log DW2
318	D[8:2]F0x178: Header Log DW3
318	D[8:2]F0x17C: Root Error Command
318	D[8:2]F0x180: Root Error Status
319	D[8:2]F0x184: Error Source ID
320	D18F0x00: Device/Vendor ID
320	D18F0x04: Status/Command
320	D18F0x08: Class Code/Revision ID
320	D18F0x0C: Header Type
320	D18F0x34: Capabilities Pointer

320	D18F0x40: Routing Table
321	D18F0x60: Node ID
321	D18F0x64: Unit ID
321	D18F0x68: Link Transaction Control
323	D18F0x6C: Link Initialization Control
323	D18F0x84: Link Control
324	D18F0x90: Upstream Base Channel Buffer Count
325	D18F0x94: Link Isochronous Channel Buffer Count
326	D18F0x98: Link Type
326	D18F0x9C: Link Frequency Extension
326	D18F0x110: Link Clumping Enable
326	D18F0x16C: Link Global Extended Control
326	D18F0x170: Link Extended Control
327	D18F0x1A0: Link Initialization Status
327	D18F0x1DC: Core Enable
328	D18F1x00: Device/Vendor ID
328	D18F1x08: Class Code/Revision ID
328	D18F1x0C: Header Type
328	D18F1x[144:140,44:40]: DRAM Base/Limit
329	D18F1x[1CC:180,BC:80]: MMIO Base/Limit
333	D18F1x[DC:C0]: IO-Space Base/Limit
335	D18F1x[EC:E0]: Configuration Map
335	D18F1xF0: DRAM Hole Address
336	D18F1xF4: VGA Enable
336	D18F1x10C: DCT Configuration Select
337	D18F1x120: DRAM Base System Address
337	D18F1x124: DRAM Limit System Address
338	D18F2x00: Device/Vendor ID
338	D18F2x08: Class Code/Revision ID
338	D18F2x0C: Header Type
338	D18F2x[5C:40]_dct[1:0]: DRAM CS Base Address
340	D18F2x[6C:60]_dct[1:0]: DRAM CS Mask
340	D18F2x78_dct[1:0]: DRAM Control
340	D18F2x7C_dct[1:0]: DRAM Initialization
341	D18F2x80_dct[1:0]: DRAM Bank Address Mapping
342	D18F2x84_dct[1:0]: DRAM MRS
343	D18F2x88_dct[1:0]: DRAM Timing Low
344	D18F2x8C_dct[1:0]: DRAM Timing High
344	D18F2x90_dct[1:0]: DRAM Configuration Low
345	D18F2x94_dct[1:0]: DRAM Configuration High
348	D18F2x98_dct[1:0]: DRAM Controller Additional Data Offset
350	D18F2x9C_dct[1:0]: DRAM Controller Additional Data Port
350	D18F2x9C_x0000_0000_dct[1:0]_mp[1:0]: DRAM Output Driver Compensation Control
351	D18F2x9C_x0000_0[3:0]0[2:1]_dct[1:0]_mp[1:0]: DRAM Write Data Timing
352	D18F2x9C_x0000_0004_dct[1:0]_mp[1:0]: DRAM Address/Command Timing Control
354	D18F2x9C_x0000_0[3:0]0[6:5]_dct[1:0]_mp[1:0]: DRAM Read DQS Timing
355	D18F2x9C_x0000_0008_dct[1:0]_mp[1:0]: DRAM Phy Control
356	D18F2x9C_x0000_000B_dct[1:0]: DRAM Phy Status Register
356	D18F2x9C_x0000_000C_dct[1:0]: DRAM Phy Miscellaneous
357	D18F2x9C_x0000_000D_dct[1:0]_mp[1:0]: DRAM Phy DLL Control
358	D18F2x9C_x0000_00[2A:10]_dct[1:0]_mp[1:0]: DRAM DQS Receiver Enable Timing
359	D18F2x9C_x0000_00[4A:30]_dct[1:0]_mp[1:0]: DRAM DQS Write Timing
361	D18F2x9C_x0000_00[51:50]_dct[1:0]: DRAM Phase Recovery Control
362	D18F2x9C_x0D0F_0[F,7:0]0[8,0]_dct[1:0]_mp[1:0]: Data Byte Pad Configuration
363	D18F2x9C_x0D0F_0[F,7:0]02_dct[1:0]: Data Byte Transmit PreDriver Calibration
364	D18F2x9C_x0D0F_0[F,7:0]04_dct[1:0]_mp[1:0]: Data Byte Pad Configuration
365	D18F2x9C_x0D0F_0[F,7:0]0[A,6]_dct[1:0]: Data Byte Transmit PreDriver Calibration 2
365	D18F2x9C_x0D0F_0[F,7:0]0F_dct[1:0]: Data Byte DLL Clock Enable
366	D18F2x9C_x0D0F_0[F,7:0]10_dct[1:0]_mp[1:0]: Data Byte DLL Power Management
367	D18F2x9C_x0D0F_0[F,7:0]11_dct[1:0]_mp[1:0]: Data Byte DLL Power Management
368	D18F2x9C_x0D0F_0[F,7:0]13_dct[1:0]_mp[1:0]: Data Byte DLL Configuration
369	D18F2x9C_x0D0F_0[F,7:0]1C_dct[1:0]_mp[1:0]: Data Byte DLL Power Management
369	D18F2x9C_x0D0F_0[F,7:0]1F_dct[1:0]_mp[1:0]: Data Byte Receiver Configuration

370 D18F2x9C_x0D0F_0[F,7:0]30_dct[1:0]: Data Byte DLL Configuration and Power Down
 371 D18F2x9C_x0D0F_0[F,7:0]31_dct[1:0]: Data Byte Fence2 Threshold
 371 D18F2x9C_x0D0F_2[F,2:0]00_dct[1:0]_mp[1:0]: Clock Pad Configuration
 372 D18F2x9C_x0D0F_2[F,2:0]02_dct[1:0]: Clock Transmit PreDriver Calibration
 372 D18F2x9C_x0D0F_[C,8,2][2:0]1F: Receiver Configuration
 373 D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0]: Clock DLL Configuration and Power Down
 374 D18F2x9C_x0D0F_4003_dct[1:0]: DCT0 M1 Fence Value
 374 D18F2x9C_x0D0F_4004_dct[1:0]: DCT1 M1 Fence Value
 374 D18F2x9C_x0D0F_4009_dct[1:0]: Phy Cmp Configuration
 376 D18F2x9C_x0D0F_8[1:0]0[8,4,0]_dct[1:0]_mp[1:0]: CAD1 Pad Configuration
 376 D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0]: Transmit PreDriver Calibration
 377 D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0]: Transmit PreDriver Calibration 2
 377 D18F2x9C_x0D0F_8[1:0]20_dct[1:0]_mp[1:0]: DLL Delay and Configuration
 378 D18F2x9C_x0D0F_8021_dct[1:0]_mp[1:0]: DLL CS 6 & 7 Timing Control
 378 D18F2x9C_x0D0F_812F_dct[1:0]: Tristate Configuration
 379 D18F2x9C_x0D0F_[C,8][F:0]30_dct[1:0]: Cmd/Addr DLL Configuration and Power Down
 379 D18F2x9C_x0D0F_C000_dct[1:0]_mp[1:0]: CKE 2.0X Pad Configuration
 380 D18F2x9C_x0D0F_C0[10,0C,08,04]_dct[1:0]_mp[1:0]: CAD2 Pad Configuration
 380 D18F2x9C_x0D0F_C020_dct[1:0]_mp[1:0]: CAD2 Address DLL Delay and Configuration
 380 D18F2x9C_x0D0F_C021_dct[1:0]_mp[1:0]: CAD2 CKE DLL Delay and Configuration
 381 D18F2x9C_x0D0F_E000_dct[1:0]_mp[1:0]: Phy Master Configuration
 381 D18F2x9C_x0D0F_E006_dct[1:0]: Phy PLL Lock Time
 381 D18F2x9C_x0D00_E008_dct[1:0]: Phy Master Configuration
 382 D18F2x9C_x0D04_E008_dct[1:0]: Phy Master Configuration
 382 D18F2x9C_x0D0F_E00A_dct[1:0]: Phy Dynamic Power Mode
 382 D18F2x9C_x0D0F_E010_dct[1:0]: DLL Gain Sequencer Timer 0
 383 D18F2x9C_x0D0F_E011_dct[1:0]: DLL Gain Sequencer Timer 1
 383 D18F2x9C_x0D0F_E012_dct[1:0]: DLL Gain Frequency Threshold
 383 D18F2x9C_x0D0F_E013_dct[1:0]: Phy PLL Regulator Wait Time
 384 D18F2x9C_x0D0F_E018_dct[1:0]: Fence
 384 D18F2x9C_x0D0F_E019_dct[1:0]: Fence2
 385 D18F2xA4: DRAM Controller Temperature Throttle
 386 D18F2xA8_dct[1:0]: DRAM Controller Miscellaneous 2
 388 D18F2xAC: DRAM Controller Temperature Status
 388 D18F2xF8: P-state Power Information 1
 388 D18F2x104: P-state Power Information 3
 388 D18F2x10C: Swap Interleaved Region Base/Limit
 389 D18F2x110: DRAM Controller Select Low
 391 D18F2x114: DRAM Controller Select High
 391 D18F2x118: Memory Controller Configuration Low
 392 D18F2x11C: Memory Controller Configuration High
 395 D18F2x1B0: Extended Memory Controller Configuration Low
 397 D18F2x1B4: Extended Memory Controller Configuration High Register
 398 D18F2x200_dct[1:0]_mp[1:0]: DRAM Timing 0
 399 D18F2x204_dct[1:0]_mp[1:0]: DRAM Timing 1
 400 D18F2x208_dct[1:0]: DRAM Timing 2
 400 D18F2x20C_dct[1:0]_mp[1:0]: DRAM Timing 3
 401 D18F2x210_dct[1:0]_nbp[3:0]: DRAM NB P-state
 402 D18F2x214_dct[1:0]_mp[1:0]: DRAM Timing 4
 403 D18F2x218_dct[1:0]_mp[1:0]: DRAM Timing 5
 404 D18F2x21C_dct[1:0]_mp[1:0]: DRAM Timing 6
 404 D18F2x220_dct[1:0]: DRAM Timing 7
 405 D18F2x224_dct[1:0]: DRAM Timing 8
 406 D18F2x228_dct[1:0]: DRAM Timing 9
 406 D18F2x22C_dct[1:0]_mp[1:0]: DRAM Timing 10
 406 D18F2x[234:230]_dct[1:0]: DRAM Read ODT Pattern [High:Low]
 407 D18F2x[23C:238]_dct[1:0]: DRAM Write ODT Pattern [High:Low]
 407 D18F2x240_dct[1:0]_mp[1:0]: DRAM ODT Control
 408 D18F2x244_dct[1:0]: DRAM Controller Miscellaneous 3
 408 D18F2x248_dct[1:0]_mp[1:0]: DRAM Power Management 0
 409 D18F2x24C_dct[1:0]: DRAM Power Management 1
 410 D18F2x250_dct[1:0]: DRAM Loopback and Training Control
 411 D18F2x25[8,4]_dct[1:0]: DRAM Target [B, A] Base

411 D18F2x260_dct[1:0]: DRAM Command 1
 411 D18F2x264_dct[1:0]: DRAM Status 0
 412 D18F2x268_dct[1:0]: DRAM Status 1
 412 D18F2x26C_dct[1:0]: DRAM Status 2
 413 D18F2x270_dct[1:0]: DRAM PRBS
 413 D18F2x274_dct[1:0]: DRAM DQ Mask Low
 413 D18F2x278_dct[1:0]: DRAM DQ Mask High
 414 D18F2x28C_dct[1:0]: DRAM Command 2
 414 D18F2x290_dct[1:0]: DRAM Status 3
 414 D18F2x294_dct[1:0]: DRAM Status 4
 415 D18F2x298_dct[1:0]: DRAM Status 5
 415 D18F2x2E0_dct[1:0]: Memory P-state Control and Status
 415 D18F2x2E8_dct[1:0]_mp[1:0]: MRS Buffer
 416 D18F2x2EC_dct[1:0]_mp[1:0]: MRS Buffer
 416 D18F2x2F0_dct[1:0]_mp[1:0]: DRAM Controller Misc 3
 416 D18F2x400_dct[1:0]: GMC to DCT Control 0
 417 D18F2x404_dct[1:0]: GMC to DCT Control 1
 417 D18F2x408_dct[1:0]: GMC to DCT Control 2
 417 D18F2x420_dct[1:0]: GMC to DCT FIFO Config 1
 419 D18F3x00: Device/Vendor ID
 419 D18F3x04: Status/Command
 419 D18F3x08: Class Code/Revision ID
 419 D18F3x0C: Header Type
 419 D18F3x34: Capability Pointer
 419 D18F3x40: MCA NB Control
 420 D18F3x44: MCA NB Configuration
 423 D18F3x48: MCA NB Status Low
 423 D18F3x4C: MCA NB Status High
 423 D18F3x50: MCA NB Address Low
 423 D18F3x54: MCA NB Address High
 423 D18F3x64: Hardware Thermal Control (HTC)
 424 D18F3x68: Software P-state Limit
 424 D18F3x6C: Data Buffer Count
 425 D18F3x70: SRI to XBAR Command Buffer Count
 426 D18F3x74: XBAR to SRI Command Buffer Count
 427 D18F3x78: MCT to XBAR Buffer Count
 428 D18F3x7C: Free List Buffer Count
 428 D18F3x[84:80]: ACPI Power State Control
 431 D18F3x88: NB Configuration Low
 431 D18F3x8C: NB Configuration High
 431 D18F3xA0: Power Control Miscellaneous
 432 D18F3xA4: Reported Temperature Control
 433 D18F3xA8: Pop Up and Down P-states
 434 D18F3xD4: Clock Power/Timing Control 0
 435 D18F3xD8: Clock Power/Timing Control 1
 436 D18F3xDC: Clock Power/Timing Control 2
 437 D18F3xE4: Thermtrip Status
 437 D18F3xE8: Northbridge Capabilities
 438 D18F3xF0: DEV Capability Header Register
 438 D18F3xF4: DEV Function Register
 439 D18F3xF8: DEV Data Port
 439 D18F3xF8_x4: DEV Secure Loader Control Register
 439 D18F3xFC: CPUID Family/Model
 440 D18F3x140: SRI to XCS Token Count
 441 D18F3x144: MCT to XCS Token Count
 441 D18F3x148: Link to XCS Token Count
 441 D18F3x17C: Extended Freelist Buffer Count
 442 D18F3x180: Extended NB MCA Configuration
 443 D18F3x190: Downcore Control
 443 D18F3x1A0: Core to NB Buffer Count
 444 D18F3x1CC: IBS Control
 444 D18F3x1FC: Product Information Register 1
 446 D18F4x00: Device/Vendor ID

446 D18F4x04: Status/Command
446 D18F4x08: Class Code/Revision ID
446 D18F4x0C: Header Type
446 D18F4x34: Capabilities Pointer
447 D18F4x108: TDP Limit 1
447 D18F4x10C: TDP Limit 2
447 D18F4x110: Sample and Residency Timers
447 D18F4x11[C:8]: C-state Control
450 D18F4x124: C-state Interrupt Control
450 D18F4x128: C-state Policy Control 1
451 D18F4x13C: SMU P-state Control
451 D18F4x15C: Core Performance Boost Control
452 D18F4x164: Fixed Errata
452 D18F4x16C: APM TDP Control
453 D18F5x00: Device/Vendor ID
453 D18F5x04: Status/Command
453 D18F5x08: Class Code/Revision ID
453 D18F5x0C: Header Type
453 D18F5x34: Capabilities Pointer
453 D18F5x[70,60,50,40]: Northbridge Performance Event Select Low
454 D18F5x[74,64,54,44]: Northbridge Performance Event Select High
454 D18F5x[78,68,58,48]: Northbridge Performance Event Counter Low
454 D18F5x[7C,6C,5C,4C]: Northbridge Performance Event Counter High
454 D18F5x80: Compute Unit Status
455 D18F5x84: Northbridge Capabilities 2
455 D18F5x88: Northbridge Configuration 4 (NB_CFG4)
456 D18F5xE0: Processor TDP Running Average
456 D18F5x128: Clock Power/Timing Control 3
457 D18F5x12C: Clock Power/Timing Control 4
458 D18F5x1[6C:60]: Northbridge P-state [3:0]
459 D18F5x170: Northbridge P-state Control
461 D18F5x174: Northbridge P-state Status
461 D18F5x178: Northbridge Fusion Configuration
462 D18F5x17C: Miscellaneous Voltages
463 D18F5x188: Clock Power/Timing Control 5
463 D18F5x194: Name String Address Port
463 D18F5x198: Name String Data Port
463 D18F5x198_x[B:0]: Name String Data
464 GMMx63C: CFG_VOLTAGE_CONTROL
464 GMMx640: CFG_VOLTAGE_STATUS
464 GMMx770: CG Voltage Control
464 GMMx774: CG Voltage Status
465 GMMx7A0: LCLK_DEEP_SLEEP_CNTL
467 IOMMUx00: Device Table Base Address Low
467 IOMMUx04: Device Table Base Address High
467 IOMMUx08: Command Buffer Base Address Low
467 IOMMUx0C: Command Buffer Base Address High
468 IOMMUx10: Event Log Base Address Low
468 IOMMUx14: Event Log Base Address High
468 IOMMUx18: Control Low
470 IOMMUx20: Exclusion Range Base Low
470 IOMMUx24: Exclusion Range Base High
470 IOMMUx28: Exclusion Range Limit Low
471 IOMMUx2C: Exclusion Range Limit High
471 IOMMUx30: Extended Feature Low
472 IOMMUx34: Extended Feature High
472 IOMMUx38: PPR Log Base Address Low
472 IOMMUx3C: PPR Log Base Address High
472 IOMMUx40: Hardware Error Upper Low
472 IOMMUx44: Hardware Error Upper High
473 IOMMUx48: Hardware Error Lower Low
473 IOMMUx4C: Hardware Error Lower High
473 IOMMUx50: Hardware Error Status

473 IOMMUx2000: Command Buffer Head Pointer
 473 IOMMUx2008: Command Buffer Tail Pointer
 474 IOMMUx2010: Event Log Head Pointer
 474 IOMMUx2018: Event Log Tail Pointer
 474 IOMMUx2020: Status
 475 IOMMUx2030: PPR Log Head Pointer
 475 IOMMUx2038: PPR Log Tail Pointer
 476 IOMMUx4000: Counter Configuration
 476 IOMMUx4008: Counter PASID Bank Lock Low
 476 IOMMUx400C: Counter PASID Bank Lock High
 477 IOMMUx4010: Domain Bank Lock Low
 477 IOMMUx4014: Domain Bank Lock High
 477 IOMMUx4018: DeviceID Bank Lock Low
 477 IOMMUx401C: DeviceID Bank Lock High
 477 IOMMUx4[1,0][3:0]00: Counter Low
 478 IOMMUx4[1,0][3:0]04: Counter High
 478 IOMMUx4[1,0][3:0]08: Counter Source
 478 IOMMUx4[1,0][3:0]10: PASID Match Low
 479 IOMMUx4[1,0][3:0]14: PASID Match High
 479 IOMMUx4[1,0][3:0]18: Domain Match Low
 480 IOMMUx4[1,0][3:0]1C: Domain Match High
 480 IOMMUx4[1,0][3:0]20: DeviceID Match Low
 480 IOMMUx4[1,0][3:0]24: DeviceID Match High
 481 IOMMUx4[1,0][3:0]28: Counter Report Low
 481 IOMMUx4[1,0][3:0]2C: Counter Report High
 482 APIC20: APIC ID
 482 APIC30: APIC Version
 482 APIC80: Task Priority (TPR)
 482 APIC90: Arbitration Priority (APR)
 483 APICA0: Processor Priority (PPR)
 483 APICB0: End of Interrupt
 483 APICC0: Remote Read
 483 APICD0: Logical Destination (LDR)
 483 APICE0: Destination Format
 484 APICF0: Spurious-Interrupt Vector (SVR)
 484 APIC[170:100]: In-Service (ISR)
 484 APIC[1F0:180]: Trigger Mode (TMR)
 485 APIC[270:200]: Interrupt Request (IRR)
 485 APIC280: Error Status
 486 APIC300: Interrupt Command Low (ICR Low)
 487 APIC310: Interrupt Command High (ICR High)
 487 APIC320: LVT Timer
 488 APIC330: LVT Thermal Sensor
 488 APIC340: LVT Performance Monitor
 488 APIC3[60:50]: LVT LINT[1:0]
 489 APIC370: LVT Error
 489 APIC380: Timer Initial Count
 489 APIC390: Timer Current Count
 489 APIC3E0: Timer Divide Configuration
 490 APIC400: Extended APIC Feature
 490 APIC410: Extended APIC Control
 491 APIC420: Specific End Of Interrupt
 491 APIC[4F0:480]: Interrupt Enable
 491 APIC[530:500]: Extended Interrupt [3:0] Local Vector Table
 493 CPUID Fn0000_0000_EAX: Processor Vendor and Largest Standard Function Number
 493 CPUID Fn0000_0000_E[D,C,B]X: Processor Vendor
 493 CPUID Fn0000_0001_EAX: Family, Model, Stepping Identifiers
 494 CPUID Fn0000_0001_EBX: LocalApicId, LogicalProcessorCount, CLFlush
 494 CPUID Fn0000_0001_ECX: Feature Identifiers
 495 CPUID Fn0000_0001_EDX: Feature Identifiers
 496 CPUID Fn0000_000[4,3,2]: Reserved
 496 CPUID Fn0000_0005_EAX: Monitor/MWait
 496 CPUID Fn0000_0005_EBX: Monitor/MWait

496 CUID Fn0000_0005_ECX: Monitor/MWait
 496 CUID Fn0000_0005_EDX: Monitor/MWait
 496 CUID Fn0000_0006_EAX: Thermal and Power Management
 496 CUID Fn0000_0006_EBX: Thermal and Power Management
 497 CUID Fn0000_0006_ECX: Thermal and Power Management
 497 CUID Fn0000_0006_EDX: Thermal and Power Management
 497 CUID Fn0000_0007_EAX_x0: Structured Extended Feature Identifiers (ECX=0)
 497 CUID Fn0000_0007_EBX_x0: Structured Extended Feature Identifiers (ECX=0)
 497 CUID Fn0000_0007_ECX_x0: Structured Extended Feature Identifiers (ECX=0)
 497 CUID Fn0000_0007_EDX_x0: Structured Extended Feature Identifiers (ECX=0)
 497 CUID Fn0000_000[A:8]: Reserved
 498 CUID Fn0000_000B_EAX: Reserved
 498 CUID Fn0000_000B_EBX: Extended Topology Enumeration
 498 CUID Fn0000_000B_ECX: Reserved
 498 CUID Fn0000_000B_EDX: Reserved
 498 CUID Fn0000_000C: Reserved
 498 CUID Fn0000_000D_EAX_x0: Processor Extended State Enumeration (ECX=0)
 498 CUID Fn0000_000D_EBX_x0: Processor Extended State Enumeration (ECX=0)
 499 CUID Fn0000_000D_ECX_x0: Processor Extended State Enumeration (ECX=0)
 499 CUID Fn0000_000D_EDX_x0: Processor Extended State Enumeration (ECX=0)
 499 CUID Fn0000_000D_EAX_x2: Processor Extended State Enumeration (ECX=2)
 499 CUID Fn0000_000D_EBX_x2: Processor Extended State Enumeration (ECX=2)
 499 CUID Fn0000_000D_ECX_x2: Processor Extended State Enumeration (ECX=2)
 499 CUID Fn0000_000D_EDX_x2: Processor Extended State Enumeration (ECX=2)
 499 CUID Fn0000_000D_EAX_x3E: Processor Extended State Enumeration (ECX=62)
 500 CUID Fn0000_000D_EBX_x3E: Processor Extended State Enumeration (ECX=62)
 500 CUID Fn0000_000D_ECX_x3E: Processor Extended State Enumeration (ECX=62)
 500 CUID Fn0000_000D_EDX_x3E: Processor Extended State Enumeration (ECX=62)
 500 CUID Fn8000_0000_EAX: Largest Extended Function Number
 500 CUID Fn8000_0000_E[D,C,B]X: Processor Vendor
 500 CUID Fn8000_0001_EAX: Family, Model, Stepping Identifiers
 501 CUID Fn8000_0001_EBX: BrandId Identifier
 501 CUID Fn8000_0001_ECX: Feature Identifiers
 502 CUID Fn8000_0001_EDX: Feature Identifiers
 503 CUID Fn8000_000[4:2]_E[D,C,B,A]X: Processor Name String Identifier
 503 CUID Fn8000_0005_EAX: L1 TLB 2M/4M Identifiers
 504 CUID Fn8000_0005_EBX: L1 TLB 4K Identifiers
 504 CUID Fn8000_0005_ECX: L1 Data Cache Identifiers
 504 CUID Fn8000_0005_EDX: L1 Instruction Cache Identifiers
 504 CUID Fn8000_0006_EAX: L2 TLB 2M/4M Identifiers
 505 CUID Fn8000_0006_EBX: L2 TLB 4K Identifiers
 505 CUID Fn8000_0006_ECX: L2 Cache Identifiers
 506 CUID Fn8000_0006_EDX: L3 Cache Identifiers
 506 CUID Fn8000_0007_E[C,B,A]X: Advanced Power Management Information
 506 CUID Fn8000_0007_EDX: Advanced Power Management Information
 507 CUID Fn8000_0008_EAX: Long Mode Address Size Identifiers
 507 CUID Fn8000_0008_EBX: Reserved
 507 CUID Fn8000_0008_ECX: APIC ID Size and Core Count
 507 CUID Fn8000_0008_EDX: Reserved
 508 CUID Fn8000_0009: Reserved
 508 CUID Fn8000_000A_EAX: SVM Revision
 508 CUID Fn8000_000A_EBX: SVM Revision and Feature Identification
 508 CUID Fn8000_000A_ECX: Reserved
 508 CUID Fn8000_000A_EDX: SVM Feature Identification
 509 CUID Fn8000_00[18:0B]: Reserved
 509 CUID Fn8000_0019_EAX: L1 TLB 1G Identifiers
 509 CUID Fn8000_0019_EBX: L2 TLB 1G Identifiers
 509 CUID Fn8000_0019_E[D,C]X: Reserved
 509 CUID Fn8000_001A_EAX: Performance Optimization Identifiers
 510 CUID Fn8000_001A_E[D,C,B]X: Reserved
 510 CUID Fn8000_001B_EAX: Instruction Based Sampling Identifiers
 510 CUID Fn8000_001B_E[D,C,B]X: Instruction Based Sampling Identifiers
 510 CUID Fn8000_001C_EAX: Lightweight Profiling Capabilities 0

511 CPUID Fn8000_001C_EBX: Lightweight Profiling Capabilities 0
 511 CPUID Fn8000_001C_ECX: Lightweight Profiling Capabilities 0
 512 CPUID Fn8000_001C_EDX: Lightweight Profiling Capabilities 0
 512 CPUID Fn8000_001D_EAX_x0: Cache Properties
 513 CPUID Fn8000_001D_EAX_x1: Cache Properties
 513 CPUID Fn8000_001D_EAX_x2: Cache Properties
 514 CPUID Fn8000_001D_EAX_x3: Cache Properties
 514 CPUID Fn8000_001D_EBX_x0: Cache Properties
 514 CPUID Fn8000_001D_EBX_x1: Cache Properties
 515 CPUID Fn8000_001D_EBX_x2: Cache Properties
 515 CPUID Fn8000_001D_EBX_x3: Cache Properties
 515 CPUID Fn8000_001D_ECX_x0: Cache Properties
 515 CPUID Fn8000_001D_ECX_x1: Cache Properties
 516 CPUID Fn8000_001D_ECX_x2: Cache Properties
 516 CPUID Fn8000_001D_ECX_x3: Cache Properties
 516 CPUID Fn8000_001D_EDX_x0: Cache Properties
 516 CPUID Fn8000_001D_EDX_x1: Cache Properties
 516 CPUID Fn8000_001D_EDX_x2: Cache Properties
 517 CPUID Fn8000_001D_EDX_x3: Cache Properties
 517 CPUID Fn8000_001E_EAX: Extended APIC ID
 517 CPUID Fn8000_001E_EBX: Compute Unit Identifiers
 517 CPUID Fn8000_001E_ECX: Node Identifiers
 517 CPUID Fn8000_001E_EDX: Reserved
 519 MSR0000_0000: Load-Store MCA Address
 519 MSR0000_0001: Load-Store MCA Status
 519 MSR0000_0010: Time Stamp Counter (TSC)
 519 MSR0000_001B: APIC Base Address (APIC_BAR)
 519 MSR0000_002A: Cluster ID (EBL_CR_POWERON)
 520 MSR0000_00E7: Max Performance Frequency Clock Count (MPERF)
 520 MSR0000_00E8: Actual Performance Frequency Clock Count (APERF)
 520 MSR0000_00FE: MTRR Capabilities (MTRRcap)
 520 MSR0000_0174: SYSENTER CS (SYSENTER_CS)
 521 MSR0000_0175: SYSENTER ESP (SYSENTER_ESP)
 521 MSR0000_0176: SYSENTER EIP (SYSENTER_EIP)
 521 MSR0000_0179: Global Machine Check Capabilities (MCG_CAP)
 521 MSR0000_017A: Global Machine Check Status (MCG_STAT)
 521 MSR0000_017B: Global Machine Check Exception Reporting Control (MCG_CTL)
 522 MSR0000_01D9: Debug Control (DBG_CTL_MSR)
 522 MSR0000_01DB: Last Branch From IP (BR_FROM)
 522 MSR0000_01DC: Last Branch To IP (BR_TO)
 522 MSR0000_01DD: Last Exception From IP
 523 MSR0000_01DE: Last Exception To IP
 523 MSR0000_020[F:0]: Variable-Size MTRRs Base/Mask
 524 MSR0000_02[6F:68,59:58,50]: Fixed-Size MTRRs
 526 MSR0000_0277: Page Attribute Table (PAT)
 527 MSR0000_02FF: MTRR Default Memory Type (MTRRdefType)
 527 MSR0000_0400: LS Machine Check Control (MC0_CTL)
 528 MSR0000_0401: LS Machine Check Status (MC0_STATUS)
 530 MSR0000_0402: LS Machine Check Address (MC0_ADDR)
 531 MSR0000_0403: LS Machine Check Miscellaneous (MC0_MISC)
 531 MSR0000_0404: IF Machine Check Control (MC1_CTL)
 532 MSR0000_0405: IF Machine Check Status (MC1_STATUS)
 537 MSR0000_0406: IF Machine Check Address (MC1_ADDR)
 539 MSR0000_0407: IF Machine Check Miscellaneous (MC1_MISC)
 539 MSR0000_0408: CU Machine Check Control (MC2_CTL)
 539 MSR0000_0409: CU Machine Check Status (MC2_STATUS)
 544 MSR0000_040A: CU Machine Check Address (MC2_ADDR)
 546 MSR0000_040B: CU Machine Check Miscellaneous (MC2_MISC)
 546 MSR0000_040C: MC3 Machine Check Control (MC3_CTL)
 546 MSR0000_040D: MC3 Machine Check Status (MC3_STATUS)
 547 MSR0000_040E: MC3 Machine Check Address (MC3_ADDR)
 547 MSR0000_040F: MC3 Machine Check Miscellaneous (MC3_MISC)
 547 MSR0000_0410: NB Machine Check Control (MC4_CTL)

548 MSR0000_0411: NB Machine Check Status (MC4_STATUS)
 551 MSR0000_0412: NB Machine Check Address (MC4_ADDR)
 554 MSR0000_0413: NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)
 556 MSR0000_0414: EX Machine Check Control (MC5_CTL)
 556 MSR0000_0415: EX Machine Check Status (MC5_STATUS)
 558 MSR0000_0416: EX Machine Check Address (MC5_ADDR)
 559 MSR0000_0417: EX Machine Check Miscellaneous (MC5_MISC)
 559 MSR0000_0418: FP Machine Check Control (MC6_CTL)
 560 MSR0000_0419: FP Machine Check Status (MC6_STATUS)
 561 MSR0000_041A: FP Machine Check Address (MC6_ADDR)
 561 MSR0000_041B: FP Machine Check Miscellaneous (MC6_MISC)
 562 MSRC000_0080: Extended Feature Enable (EFER)
 562 MSRC000_0081: SYSCALL Target Address (STAR)
 562 MSRC000_0082: Long Mode SYSCALL Target Address (STAR64)
 562 MSRC000_0083: Compatibility Mode SYSCALL Target Address (STARCOMPAT)
 563 MSRC000_0084: SYSCALL Flag Mask (SYSCALL_FLAG_MASK)
 563 MSRC000_00E7: Read-Only Max Performance Frequency Clock Count (MPerfReadOnly)
 563 MSRC000_00E8: Read-Only Actual Performance Frequency Clock Count (APerfReadOnly)
 563 MSRC000_0100: FS Base (FS_BASE)
 564 MSRC000_0101: GS Base (GS_BASE)
 564 MSRC000_0102: Kernel GS Base (KernelGSbase)
 564 MSRC000_0103: Auxiliary Time Stamp Counter (TSC_AUX)
 564 MSRC000_0104: Time Stamp Counter Ratio (TscRateMsr)
 564 MSRC000_0105: Lightweight Profile Configuration (LWP_CFG)
 565 MSRC000_0106: Lightweight Profile Control Block Address (LWP_CBADDR)
 565 MSRC000_0408: Machine Check Misc 4 (Link Thresholding) 1 (MC4_MISC1)
 566 MSRC000_040[F:9]: Reserved
 567 MSRC001_00[03:00]: Performance Event Select (PERF_CTL[3:0])
 567 MSRC001_00[07:04]: Performance Event Counter (PERF_CTR[3:0])
 567 MSRC001_0010: System Configuration (SYS_CFG)
 568 MSRC001_0015: Hardware Configuration (HWCR)
 570 MSRC001_00[18,16]: IO Range Base (IORR_BASE[1:0])
 570 MSRC001_00[19,17]: IO Range Mask (IORR_MASK[1:0])
 571 MSRC001_001A: Top Of Memory (TOP_MEM)
 571 MSRC001_001D: Top Of Memory 2 (TOM2)
 571 MSRC001_001F: Northbridge Configuration 1 (NB_CFG1)
 572 MSRC001_0022: Machine Check Exception Redirection
 572 MSRC001_00[35:30]: Processor Name String
 573 MSRC001_003E: Hardware Thermal Control (HTC)
 573 MSRC001_0044: LS Machine Check Control Mask (MC0_CTL_MASK)
 574 MSRC001_0045: IF Machine Check Control Mask (MC1_CTL_MASK)
 574 MSRC001_0046: CU Machine Check Control Mask (MC2_CTL_MASK)
 575 MSRC001_0047: Reserved (MC3_CTL_MASK)
 575 MSRC001_0048: NB Machine Check Control Mask (MC4_CTL_MASK)
 576 MSRC001_0049: EX Machine Check Control Mask (MC5_CTL_MASK)
 576 MSRC001_004A: FP Machine Check Control Mask (MC6_CTL_MASK)
 576 MSRC001_00[53:50]: IO Trap (SMI_ON_IO_TRAP_[3:0])
 577 MSRC001_0054: IO Trap Control (SMI_ON_IO_TRAP_CTL_STS)
 577 MSRC001_0055: Interrupt Pending
 578 MSRC001_0056: SMI Trigger IO Cycle
 578 MSRC001_0058: MMIO Configuration Base Address
 579 MSRC001_0061: P-state Current Limit
 579 MSRC001_0062: P-state Control
 579 MSRC001_0063: P-state Status
 580 MSRC001_00[6B:64]: P-state [7:0]
 581 MSRC001_0070: COFVID Control
 582 MSRC001_0071: COFVID Status
 583 MSRC001_0073: C-state Base Address
 583 MSRC001_0111: SMM Base Address (SMM_BASE)
 583 MSRC001_0112: SMM TSeg Base Address (SMMAddr)
 584 MSRC001_0113: SMM TSeg Mask (SMMMask)
 585 MSRC001_0114: Virtual Machine Control (VM_CR)
 586 MSRC001_0115: IGNNE

586 MSRC001_0116: SMM Control (SMM_CTL)
 586 MSRC001_0117: Virtual Machine Host Save Physical Address (VM_HSAVE_PA)
 586 MSRC001_0118: SVM Lock Key
 587 MSRC001_011A: Local SMI Status
 587 MSRC001_0140: OS Visible Work-around MSR0 (OSVW_ID_Length)
 587 MSRC001_0141: OS Visible Work-around MSR1 (OSVW_Status)
 587 MSRC001_020[A,8,6,4,2,0]: Performance Event Select (PERF_CTL[5:0])
 589 MSRC001_020[B,9,7,5,3,1]: Performance Event Counter (PERF_CTR[5:0])
 589 MSRC001_024[6,4,2,0]: Northbridge Performance Event Select (NB_PERF_CTL[3:0])
 590 MSRC001_024[7,5,3,1]: Northbridge Performance Event Counter (NB_PERF_CTR[3:0])
 592 MSRC001_1002: CPUID Features for CPUID Fn0000_0007_E[B,A]X_x0
 592 MSRC001_1003: Thermal and Power Management CPUID Features
 592 MSRC001_1004: CPUID Features (Features)
 593 MSRC001_1005: Extended CPUID Features (ExtFeatures)
 595 MSRC001_100C: Node ID
 596 MSRC001_1020: Load-Store Configuration (LS_CFG)
 596 MSRC001_1021: Instruction Cache Configuration (IC_CFG)
 596 MSRC001_1022: Data Cache Configuration (DC_CFG)
 596 MSRC001_1023: Combined Unit Configuration (CU_CFG)
 597 MSRC001_1028: Floating Point Configuration (FP_CFG)
 597 MSRC001_102A: Combined Unit Configuration 2 (CU_CFG2)
 598 MSRC001_102B: Combined Unit Configuration 3 (CU_CFG3)
 599 MSRC001_102D: Load-Store Configuration 2 (LS_CFG2)
 599 MSRC001_1030: IBS Fetch Control (IC_IBS_CTL)
 600 MSRC001_1031: IBS Fetch Linear Address (IC_IBS_LIN_AD)
 600 MSRC001_1032: IBS Fetch Physical Address (IC_IBS_PHYS_AD)
 601 MSRC001_1033: IBS Execution Control (SC_IBS_CTL)
 602 MSRC001_1034: IBS Op Logical Address (IBSOP_RIP)
 602 MSRC001_1035: IBS Op Data (SC_IBS_DATA)
 602 MSRC001_1036: IBS Op Data 2 (IbsOpData2)
 603 MSRC001_1037: IBS Op Data 3 (DC_IBS_DATA, IbsOpData3)
 604 MSRC001_1038: IBS DC Linear Address (DC_IBS_LIN_ADDR)
 604 MSRC001_1039: IBS DC Physical Address (DC_IBS_PHYS_ADDR)
 604 MSRC001_103A: IBS Control
 604 MSRC001_103B: IBS Branch Target Address (BP_IBSTGT_RIP)
 606 PMCx000: FPU Pipe Assignment
 606 PMCx001: FP Scheduler Empty
 606 PMCx003: Retired SSE/BNI Ops
 607 PMCx004: Number of Move Elimination and Scalar Op Optimization
 607 PMCx005: Retired Serializing Ops
 607 PMCx006: Number of Cycles that a Bottom-Execute uop is in the FP Scheduler
 607 PMCx020: Segment Register Loads
 608 PMCx021: Pipeline Restart Due to Self-Modifying Code
 608 PMCx022: Pipeline Restart Due to Probe Hit
 608 PMCx023: Load Queue/Store Queue Full
 608 PMCx024: Locked Operations
 608 PMCx026: Retired CLFLUSH Instructions
 609 PMCx027: Retired CPUID Instructions
 609 PMCx029: LS Dispatch
 609 PMCx02A: Canceled Store to Load Forward Operations
 609 PMCx02B: SMIs Received
 609 PMCx030: Executed CLFLUSH Instructions
 609 PMCx032: Misaligned Stores
 609 PMCx034: FP Load Buffer Stall
 609 PMCx040: Data Cache Accesses
 610 PMCx041: Data Cache Misses
 610 PMCx042: Data Cache Refills from L2 or System
 610 PMCx043: Data Cache Refills from System
 610 PMCx045: Unified TLB Hit
 611 PMCx046: Unified TLB Miss
 611 PMCx047: Misaligned Accesses
 611 PMCx04B: Prefetch Instructions Dispatched
 611 PMCx052: Ineffective Software Prefetches

611 PMCx054: Global TLB Flushes
612 PMCx065: Memory Requests by Type
612 PMCx067: Data Prefetcher
612 PMCx068: MAB Requests
612 PMCx069: MAB Wait Cycles
613 PMCx06C: Response From System on Cache Refills
613 PMCx06D: Octwords Written to System
613 PMCx076: CPU Clocks not Halted
613 PMCx07D: Requests to L2 Cache
614 PMCx07E: L2 Cache Misses
614 PMCx07F: L2 Fill/Writeback
614 PMCx165: Page Splintering
615 PMCx16C: L2 Prefetcher Trigger Events
615 PMCx080: Instruction Cache Fetches
615 PMCx081: Instruction Cache Misses
615 PMCx082: Instruction Cache Refills from L2
615 PMCx083: Instruction Cache Refills from System
615 PMCx084: L1 ITLB Miss, L2 ITLB Hit
616 PMCx085: L1 ITLB Miss, L2 ITLB Miss
616 PMCx086: Pipeline Restart Due to Instruction Stream Probe
616 PMCx087: Instruction Fetch Stall
616 PMCx088: Return Stack Hits
616 PMCx089: Return Stack Overflows
616 PMCx08B: Instruction Cache Victims
617 PMCx08C: Instruction Cache Lines Invalidated
617 PMCx099: ITLB Reloads
617 PMCx09A: ITLB Reloads Aborted
617 PMCx0C0: Retired Instructions
617 PMCx0C1: Retired uops
617 PMCx0C2: Retired Branch Instructions
617 PMCx0C3: Retired Mispredicted Branch Instructions
617 PMCx0C4: Retired Taken Branch Instructions
618 PMCx0C5: Retired Taken Branch Instructions Mispredicted
618 PMCx0C6: Retired Far Control Transfers
618 PMCx0C7: Retired Branch Resyncs
618 PMCx0C8: Retired Near Returns
618 PMCx0C9: Retired Near Returns Mispredicted
618 PMCx0CA: Retired Indirect Branches Mispredicted
618 PMCx0CB: Retired MMX/FP Instructions
618 PMCx0CD: Interrupts-Masked Cycles
619 PMCx0CE: Interrupts-Masked Cycles with Interrupt Pending
619 PMCx0CF: Interrupts Taken
619 PMCx0D0: Decoder Empty
619 PMCx0D1: Dispatch Stalls
619 PMCx0D3: Microsequencer Stall due to Serialization
619 PMCx0D5: Dispatch Stall for Instruction Retire Q Full
619 PMCx0D6: Dispatch Stall for Integer Scheduler Queue Full
620 PMCx0D7: Dispatch Stall for FP Scheduler Queue Full
620 PMCx0D8: Dispatch Stall for LDQ Full
620 PMCx0D9: Microsequencer Stall Waiting for All Quiet
620 PMCx0DB: FPU Exceptions
620 PMCx0D[F:C]: DR[3:0] Breakpoint Matches
621 PMCx1C0: Retired x87 Floating Point Operations
621 PMCx1CF: Tagged IBS Ops
621 PMCx1D8: Dispatch Stall for STQ Full
622 NBPMCx0E0: DRAM Accesses
622 NBPMCx0E1: DRAM Controller Page Table Overflows
623 NBPMCx0E2: Memory Controller DRAM Command Slots Missed
623 NBPMCx0E3: Memory Controller Turnarounds
623 NBPMCx0E4: Memory Controller Bypass Counter Saturation
624 NBPMCx0E8: Thermal Status
624 NBPMCx0E9: CPU/IO Requests to Memory/IO
625 NBPMCx0EA: Cache Block Commands

- 625 NBPMCx0EB: Sized Commands
- 625 NBPMCx0EC: Probe Responses and Upstream Requests
- 626 NBPMCx1E0: CPU to DRAM Requests to Target Node
- 627 NBPMCx1E1: IO to DRAM Requests to Target Node
- 627 NBPMCx1E2: CPU Read Command Latency to Target Node 0-3
- 627 NBPMCx1E3: CPU Read Command Requests to Target Node 0-3
- 628 NBPMCx1E4: CPU Read Command Latency to Target Node 4-7
- 628 NBPMCx1E5: CPU Read Command Requests to Target Node 4-7
- 628 NBPMCx1E6: CPU Command Latency to Target Node 0-3/4-7
- 629 NBPMCx1E7: CPU Requests to Target Node 0-3/4-7
- 629 NBPMCx1EB: Request Cache Status 1
- 630 NBPMCx1F0: Memory Controller Requests