

**BIOS and Kernel
Developer's Guide
(BKDG)
for AMD Family 15h
Models 00h-0Fh
Processors**

© 2013 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to discontinue or make changes to products, specifications, product descriptions or documentation at any time without notice. The information contained herein may be of a preliminary or advance nature. No license, whether express, implied, arising by estoppel, or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD Arrow logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Dolby is a trademark of Dolby Laboratories.

HDMI is a trademark of HDMI Licensing, LLC.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

Microsoft, Windows, Windows Vista, and DirectX are registered trademarks of Microsoft Corporation.

MMX is a trademark of Intel Corporation.

OpenCL is a trademark of Apple Inc. used by permission by Khronos.

PCIe is a registered trademark of PCI-Special Interest Group (PCI-SIG).

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

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-2 STANDARD IS EXPRESSLY PROHIBITED WITHOUT A LICENSE UNDER APPLICABLE PATENTS IN THE MPEG-2 PATENT PORTFOLIO, WHICH LICENSE IS AVAILABLE FROM MPEG LA, L.L.C., 6312 S. FIDDLERS GREEN CIRCLE, SUITE 400E, GREENWOOD VILLAGE, COLORADO 80111.

Disclaimer

While every precaution has been taken in the preparation of this document, Advanced Micro Devices, Inc. assumes no liability with respect to the operation or use of AMD hardware, software or other products and documentation described herein, for any act or omission of AMD concerning such products or this documentation, for any interruption of service, loss or interruption of business, loss of anticipatory profits, or for punitive, incidental or consequential damages in connection with the furnishing, performance, or use of the AMD hardware, software, or other products and documentation provided herein. Ensure that you have the latest documentation.

Table of Contents

1	Overview	20
1.1	Intended Audience	20
1.2	Reference Documents	20
1.3	Conventions	20
1.3.1	Numbering	20
1.3.2	Arithmetic And Logical Operators	21
1.4	Definitions	22
1.5	Changes Between Revisions and Product Variations	27
1.5.1	Revision Conventions	27
1.5.2	Major Changes Relative to Family 10h Revision D Processors	28
1.5.2.1	Major Changes to Performance Counters Relative to Fam10h RevD	32
1.5.3	Changes For Revision C	33
1.5.3.1	Performance monitor changes for revision C	33
2	Functional Description	34
2.1	Processor Overview	34
2.2	System Overview	35
2.3	Processor Initialization	36
2.3.1	BSP initialization	36
2.3.2	AP initialization	37
2.3.3	Using L2 Cache as General Storage During Boot	37
2.3.3.1	Using L2 Cache as General Storage During Boot for (PROC<OR_C0)	39
2.3.4	Multiprocessing Capability Detection	39
2.3.5	SLIT and SRAT	40
2.3.5.1	SLIT	40
2.3.5.2	SRAT	41
2.4	Processor Core	41
2.4.1	Compute Unit	41
2.4.1.1	Registers Shared by Cores in a Compute Unit	41
2.4.2	Virtual Address Space	43
2.4.3	Processor Cores and Downcoring	43
2.4.3.1	Software Downcoring using D18F3x190[DisCore]	43
2.4.4	Physical Address Space	44
2.4.5	System Address Map	44
2.4.5.1	Memory Access to the Physical Address Space	44
2.4.5.1.1	Determining Memory Type	45
2.4.5.1.2	Determining The Access Destination for Core Accesses	45
2.4.6	Timers	45
2.4.7	Implicit Conditions for TLB Invalidation	45
2.4.8	Interrupts	46
2.4.8.1	Local APIC	46
2.4.8.1.1	Detecting and Enabling	46
2.4.8.1.2	APIC Register Space	46
2.4.8.1.3	ApicId Enumeration Requirements	47
2.4.8.1.4	Physical Destination Mode	47
2.4.8.1.5	Logical Destination Mode	47
2.4.8.1.6	Interrupt Delivery	48
2.4.8.1.7	Vectored Interrupt Handling	48

2.4.8.1.8	Interrupt Masking	48
2.4.8.1.9	Spurious Interrupts	48
2.4.8.1.10	Spurious Interrupts Caused by Timer Tick Interrupt	49
2.4.8.1.11	Lowest-Priority Interrupt Arbitration	49
2.4.8.1.12	Inter-Processor Interrupts	49
2.4.8.1.13	APIC Timer Operation	50
2.4.8.1.14	Generalized Local Vector Table	50
2.4.8.1.15	State at Reset	50
2.4.8.2	System Management Mode (SMM)	50
2.4.8.2.1	SMM Overview	50
2.4.8.2.2	Operating Mode and Default Register Values	51
2.4.8.2.3	SMI Sources And Delivery	51
2.4.8.2.4	SMM Initial State	51
2.4.8.2.5	SMM Save State	52
2.4.8.2.6	Exceptions and Interrupts in SMM	58
2.4.8.2.7	The Protected ASeg and TSeg Areas	58
2.4.8.2.8	SMM Special Cycles	58
2.4.8.2.9	Locking SMM	58
2.4.8.2.10	Synchronizing SMM Entry (Spring-Boarding)	58
2.4.9	Secure Virtual Machine Mode (SVM)	60
2.4.9.1	BIOS support for SVM Disable	60
2.4.10	CPUID Instruction	60
2.4.10.1	Multi-Core Support	60
2.4.10.2	L3 Cache Support	61
2.5	Power Management	61
2.5.1	Processor Power Planes And Voltage Control	61
2.5.1.1	Internal VID Registers	62
2.5.1.1.1	VID Encodings	62
2.5.1.1.2	MinVid and MaxVid Check	62
2.5.1.2	Serial VID Interface	62
2.5.1.3	BIOS Requirements for Power Plane Initialization	62
2.5.1.4	Low Power Features	63
2.5.1.4.1	PSI_L Bit	63
2.5.1.5	Voltage Transitions	63
2.5.1.5.1	Hardware-Initiated Voltage Transitions	64
2.5.1.5.2	Software-Initiated Voltage Transitions	64
2.5.1.5.2.1	Software-Initiated NB Voltage Transitions	64
2.5.1.5.2.2	Software-Initiated Core Voltage Transitions	64
2.5.2	P-states	64
2.5.2.1	Core P-states	65
2.5.2.1.1	Application Power Management (APM)	65
2.5.2.1.1.1	TDP Limiting	66
2.5.2.1.1.2	Notification of TDP Limit Changes	66
2.5.2.1.2	Core P-state Naming and Numbering	67
2.5.2.1.2.1	Software P-state Numbering	67
2.5.2.1.2.2	Hardware P-state Numbering	68
2.5.2.1.3	Core P-state Control	68
2.5.2.1.4	Core P-state Visibility	69
2.5.2.1.5	Core P-state Limits	69
2.5.2.1.6	Core P-state Bandwidth Requirements	69
2.5.2.1.7	Core P-state Transition Behavior	70

2.5.2.1.8	BIOS Requirements for Core P-state Initialization and Transitions . . .	71
2.5.2.1.9	Processor-Systemboard Power Delivery Compatibility Check	71
2.5.2.1.10	BIOS COF and VID Requirements After Warm Reset	72
2.5.2.1.10.1	Core Maximum P-state Transition Sequence After Warm Reset	73
2.5.2.1.10.2	Core Minimum P-state Transition Sequence After Warm Reset	73
2.5.2.1.10.3	NB COF and VID Transition Sequence After Warm Reset	73
2.5.2.1.11	ACPI Processor P-state Objects	74
2.5.2.1.11.1	_PCT (Performance Control)	74
2.5.2.1.11.2	_PSS (Performance Supported States)	74
2.5.2.1.11.3	_PPC (Performance Present Capabilities)	75
2.5.2.1.11.4	_PSD (P-state Dependency)	75
2.5.2.1.12	XPSS (Microsoft® Extended PSS) Object	75
2.5.2.2	NB P-states	76
2.5.2.2.1	NB P-state Control	76
2.5.2.3	Mixed Frequency and Power P-state Configuration	76
2.5.2.3.1	Mixed Northbridge Frequency Configuration Sequence	77
2.5.2.3.2	Mixed Power P-State Configuration Sequence	77
2.5.2.3.3	Mixed Frequency and Power P-State Configuration Rules	78
2.5.2.3.4	Mixed Frequency and Power P-State Configuration Sequence	78
2.5.3	C-states	84
2.5.3.1	C-state Names and Numbers	84
2.5.3.2	C-state Request Interface	84
2.5.3.3	C-state Actions	84
2.5.3.3.1	C-state Divisors and Probes	84
2.5.3.3.2	C-state Cache Flush	84
2.5.3.3.3	Core C6 (CC6) State	85
2.5.3.4	Exiting C-states	85
2.5.3.5	C1 Enhanced State (C1E)	86
2.5.3.5.1	Message Triggered C1E	86
2.5.3.5.2	BIOS Requirements to Initialize Message Triggered C1E	87
2.5.3.6	ACPI Processor C-state Objects	87
2.5.3.6.1	_CST	87
2.5.3.6.2	_CSD	87
2.5.3.6.3	_CRS	88
2.5.3.6.4	Fixed ACPI Description Table (FADT) Entries	88
2.5.3.7	BIOS Requirements for Initialization	88
2.5.4	Frequency and Voltage Domain Dependencies	88
2.5.4.1	Dependencies Between Cores in a Compute Unit	88
2.5.5	S-states	89
2.5.5.1	ACPI Suspend to RAM State (S3)	89
2.5.6	Effective Frequency Interface	89
2.6	Processor State Transition Sequences	90
2.6.1	ACPI Power State Transitions	90
2.7	Performance Monitoring	90
2.7.1	Core Performance Monitor Counters	91
2.7.2	NB Performance Monitor Counters	92
2.7.3	Instruction Based Sampling (IBS)	93
2.8	Configuration Space	93
2.8.1	MMIO Configuration Coding Requirements	94
2.8.2	MMIO Configuration Ordering	94
2.8.3	Processor Configuration Space	94

2.9	Northbridge (NB)	95
2.9.1	NB Architecture	95
2.9.2	GART	95
2.9.3	NB Routing	96
2.9.3.1	Address Space Routing	96
2.9.3.1.1	DRAM and MMIO Memory Space	96
2.9.3.1.2	IO Space	97
2.9.3.1.3	Configuration Space	97
2.9.3.2	HyperTransport™ Technology Routing	97
2.9.3.2.1	Routing Table Configuration	99
2.9.3.2.2	NodeId Enumeration Requirements for Dual-node Processors	100
2.9.3.2.3	BIOS Requirements for Systems with Mixed Processor Families	100
2.9.3.2.4	Link Traffic Distribution	100
2.9.3.2.4.1	Coherent Link Traffic Distribution	101
2.9.3.2.4.2	Coherent Link Pair Traffic Distribution	101
2.9.3.2.4.3	Victim Distribution Mode	101
2.9.3.2.5	Display Refresh And IFCM	102
2.9.3.2.5.1	Recommended Buffer Count Settings Overview	103
2.9.4	The Level 3 Cache (L3)	104
2.9.4.1	Probe Filter	104
2.9.4.1.1	Probe Filter Errors	105
2.9.4.2	Probe Filter and ATM Mode Initialization Sequence	105
2.9.4.3	L3 Cache Partitioning	106
2.9.5	Memory Scrubbers	106
2.10	DRAM Controllers (DCTs)	106
2.10.1	DCT Configuration Registers	114
2.10.2	DDR Pad to Processor Pin Mapping	114
2.10.3	DRAM Controller Direct Response Mode	116
2.10.4	DRAM Data Burst Mapping	116
2.10.5	DCT/DRAM Initialization and Resume	116
2.10.5.1	Low Voltage DDR3	117
2.10.5.2	NB P-state Specific Configuration	117
2.10.5.3	DDR Phy Initialization	118
2.10.5.3.1	Phy Voltage Level Programming	118
2.10.5.3.2	DRAM Channel Frequency Change	118
2.10.5.3.2.1	Requirements for DRAM Frequency Change During Training	119
2.10.5.3.3	Phy Fence Programming	119
2.10.5.3.3.1	Phy Fence Training	120
2.10.5.3.4	Phy Compensation Initialization	120
2.10.5.4	SPD ROM-Based Configuration	130
2.10.5.4.1	LRDIMM Module Delay Time SPD Bytes	131
2.10.5.5	Non-SPD ROM-Based Configuration	131
2.10.5.5.1	TrdrdSdSc, TrdrdSdDc, and TrdrdDd (Read to Read Timing)	132
2.10.5.5.2	TwrwrSdSc, TwrwrSdDc, TwrwrDd (Write to Write Timing)	132
2.10.5.5.3	Twrrd (Write to Read DIMM Termination Turn-around)	133
2.10.5.5.4	TrwtTO (Read-to-Write Turnaround for Data, DQS Contention)	133
2.10.5.5.5	DRAM ODT Control	133
2.10.5.5.6	DRAM Address Timing and Output Driver Compensation Control	150
2.10.5.6	DCT Training Specific Configuration	162
2.10.5.7	DRAM Device and Controller Initialization	163
2.10.5.7.1	Software DDR3 Device Initialization	163

2.10.5.7.1.1	DDR3 MR Initialization	164
2.10.5.7.1.2	Software Control Word Initialization	167
2.10.5.8	DRAM Training	191
2.10.5.8.1	Write Levelization Training	191
2.10.5.8.1.1	Write Leveling Seed Value	193
2.10.5.8.2	DQS Receiver Enable Training	194
2.10.5.8.2.1	DQS Receiver Enable Training Seed Value	195
2.10.5.8.3	DQS Receiver Enable Cycle Training	197
2.10.5.8.4	DQS Position Training	198
2.10.5.8.5	Calculating MaxRdLatency	200
2.10.5.8.5.1	MaxRdLatency Training	201
2.10.5.8.6	Continuous Pattern Generation	201
2.10.5.8.6.1	DRAM Training Pattern Generation	201
2.10.5.9	DRAM Channel Disable	204
2.10.5.10	DRAM Phy Power Savings	204
2.10.5.10.1	Platform Specific DRAM Phy Tristate Recommendations	205
2.10.6	Memory Interleaving Modes	211
2.10.6.1	Chip Select Interleaving	212
2.10.6.2	Channel Interleaving	213
2.10.6.3	Node Interleaving	214
2.10.7	Memory Hoisting	214
2.10.7.1	DramHoleOffset Programming	215
2.10.7.2	DctSelBaseOffset Programming	215
2.10.8	DRAM CC6 Storage	216
2.10.9	On-Line Spare	217
2.10.9.1	On-Line Spare and CS Interleaving	217
2.10.10	DRAM On DIMM Thermal Management and Power Capping	218
2.10.11	LRDIMM	219
2.10.11.1	LRDIMM Rank Multiplication	219
2.11	Thermal Functions	219
2.11.1	The Tctl Temperature Scale	219
2.11.2	Temperature-Driven Logic	220
2.11.2.1	PROCHOT_L and Hardware Thermal Control (HTC)	220
2.11.2.2	Software P-state Limit Control	221
2.11.2.3	THERMTRIP	221
2.12	Links	221
2.12.1	Link Initialization	222
2.12.1.1	Ganging And Unganging	222
2.12.1.2	Ganging Detection And Control	222
2.12.1.3	Link Type Detect	222
2.12.1.3.1	Link Specific Registers	223
2.12.1.3.2	Unused Links	223
2.12.1.4	Legal Topologies	223
2.12.1.5	Link Mapping Between Package and Node	223
2.12.2	Termination and Compensation	224
2.12.3	Equalization	225
2.12.4	Link Bandwidth Requirements	225
2.12.5	Link Retry	225
2.12.6	Link LDTSTOP_L Disconnect-Reconnect	226
2.12.7	LDTSTOP_L Requirements	226
2.12.8	Response Ordering	227

2.12.9	Link Testing, BIST, and ILM	228
2.12.10	Miscellaneous Behaviors and Requirements	228
2.12.10.0.1	Legacy x87 FPU Stack Registers	229
2.13	RAS and Advanced Server Features	229
2.13.1	Machine Check Architecture	229
2.13.1.1	Machine Check Registers	229
2.13.1.2	Machine Check Errors Classes	230
2.13.1.3	Error Detection, Action, Logging, and Reporting	231
2.13.1.3.1	MCA conditions that cause Shutdown	232
2.13.1.3.2	Error Logging During Overflow	233
2.13.1.4	MCA Initialization	233
2.13.1.5	Error Code	234
2.13.1.6	Handling Machine Check Exceptions	236
2.13.1.6.1	MCA Differentiation Between System-Fatal and Process-Fatal Errors	237
2.13.1.7	Error Thresholding	238
2.13.1.8	Scrub Rate Considerations	239
2.13.1.9	Error Diagnosis	240
2.13.1.9.1	Common Diagnosis Information	241
2.13.1.10	Deferred Errors and Data Poisoning	242
2.13.1.10.1	Software Error Handling	243
2.13.2	DRAM Considerations for ECC	243
2.13.2.1	ECC Mapping	243
2.13.2.2	Single Device Data Correction (Chipkill)	244
2.13.2.3	Hardware Managed ECC History Mechanism	245
2.13.2.4	Software Managed Bad Symbol Identification	245
2.13.2.5	ECC Syndromes	246
2.13.2.5.1	x8 ECC	246
2.13.2.5.2	x4 ECC	250
2.13.3	Error Injection and Simulation	252
2.13.3.1	DRAM Error Injection	252
2.13.4	Sideband Interface (SBI)	253
2.13.4.1	SBI Processor Information	254
3	Registers	255
3.1	Register Descriptions and Mnemonics	255
3.1.1	Northbridge MSRs In Multi-Core Products	258
3.1.2	Software Recommendation (BIOS, SBIOS)	258
3.1.3	See Keyword (See:)	258
3.1.4	Mapping Tables	259
3.1.4.1	Register Mapping	259
3.1.4.2	Index Mapping	259
3.1.4.3	Field Mapping	259
3.1.4.4	Broadcast Mapping	259
3.1.4.5	Reset Mapping	259
3.1.4.6	Valid Values	259
3.1.4.7	BIOS Recommendations	260
3.2	IO Space Registers	260
3.3	Device [1F:18]h Function 0 Configuration Registers	261
3.4	Device [1F:18]h Function 1 Configuration Registers	287
3.5	Device [1F:18]h Function 2 Configuration Registers	300
3.6	Device [1F:18]h Function 3 Configuration Registers	374

3.7	Device [1F:18]h Function 4 Configuration Registers	418
3.8	Device [1F:18]h Function 5 Configuration Registers	456
3.9	APIC Registers	464
3.10	CPUID Instruction Registers	475
3.11	MSRs - MSR0000_xxxx	501
3.12	MSRs - MSRC000_0xxx	549
3.13	MSRs - MSRC001_0xxx	555
3.14	MSRs - MSRC001_1xxx	584
3.15	Core Performance Counter Events	599
3.15.1	PMCx0[1F:00] Events (FP)	599
3.15.2	PMCx0[3F:20] Events (LS)	601
3.15.3	PMCx0[5F:40] Events (DC)	603
3.15.4	PMCx[1,0][7F:60] Events (CU)	605
3.15.5	PMCx[1,0][9F:80] Events (IC)	608
3.15.6	PMCx0[BF:A0] Events	610
3.15.7	PMCx[1,0][DF:C0] Events (EX, DE)	610
3.16	NB Performance Counter Events	615
3.16.1	NBPMCx0E[7:0] Events (Memory Controller)	615
3.16.2	NBPMCx0E[F:8] Events (Crossbar)	617
3.16.3	NBPMCx0F[F:0] Events (Link, Crossbar)	620
3.16.4	NBPMCx1E[F:0] Events (Crossbar)	621
3.16.5	NBPMCx1F[F:0] Events (Memory Controller, Crossbar)	625
3.16.6	NBPMCx[5F:4E][F:0] Events (L3 Cache)	625
4	Register List	628

List of Figures

Figure 1:	A Single-Node Processor	34
Figure 2:	A Dual-Node Processor	35
Figure 3:	System Diagram	36
Figure 4:	Example 8-Node System in Twisted Ladder Topology	40
Figure 5:	Two Processor G34r1 Topology.....	98
Figure 6:	Four Processor G34r1 Maximum Performance Topology	98
Figure 7:	Four Processor G34r1 Maximum IO Topology	99
Figure 8:	Sample Four-Node Configuration.....	99
Figure 9:	Example 4-Node System in Full Topology.....	102
Figure 10:	DQS Position Training Example Results	199
Figure 11:	DQS Position Training Insertion Delay Recovery Example Results.....	200
Figure 12:	Example Cases for Programming DramHoleOffset.....	215
Figure 13:	Example Cases for Programming DctSelBaseOffset.....	216
Figure 14:	Tctl scale	220
Figure 15:	Dual-Node Processor Link Diagram.....	224
Figure 16:	Link DC Termination Mode.....	225
Figure 17:	Example of Line Interleaving from x4 DRAM Using x4 and x8 ECC.....	244
Figure 18:	Address/Command Timing at the Processor Pins	316
Figure 19:	Link Phy Recovered Clock and Sample Clock.....	429

List of Tables

Table 1:	Arithmetic and Logical Operators.....	21
Table 2:	Functions.....	21
Table 3:	Definitions.....	22
Table 4:	Processor revision conventions.....	28
Table 5:	SLIT Table Without Probe Filter.....	40
Table 6:	SLIT Table With Probe Filter.....	41
Table 7:	Compute Unit Definitions.....	41
Table 8:	SharedC and SharedNC MSR's.....	42
Table 9:	SMM Initial State.....	52
Table 10:	SMM Save State.....	52
Table 11:	Power Management Support.....	61
Table 12:	Software P-state Naming.....	67
Table 13:	Software P-state Control.....	68
Table 14:	Representative mixed frequency P-state table example (step 2).....	79
Table 15:	Representative mixed frequency P-state table example (step 3).....	79
Table 16:	Representative mixed frequency P-state table example (step 4).....	80
Table 17:	Representative mixed frequency P-state table example (step 6).....	81
Table 18:	Representative mixed frequency P-state table example (step 7).....	82
Table 19:	Representative mixed frequency P-state table example (final).....	83
Table 20:	Representative mixed frequency _PSS object example.....	83
Table 21:	Dependencies Between Cores in a Compute Unit.....	89
Table 22:	Core PMC mapping to PERF_CTL[5:0].....	91
Table 23:	Link Definitions.....	103
Table 24:	L3 Controller Definitions.....	104
Table 25:	DCT Definitions.....	107
Table 26:	DDR3 UDIMM Maximum Frequency Support for G34.....	108
Table 27:	DDR3 RDIMM Maximum Frequency Support for G34.....	109
Table 28:	DDR3 LRDIMM Maximum Frequency Support for G34.....	109
Table 29:	DDR3 UDIMM Maximum Frequency Support for C32.....	110
Table 30:	DDR3 RDIMM Maximum Frequency Support for C32.....	110
Table 31:	DDR3 LRDIMM Maximum Frequency Support for C32.....	111
Table 32:	DDR3 UDIMM Maximum Frequency Support for AM3r2 Desktop.....	112
Table 33:	DDR3 UDIMM Maximum Frequency Support for AM3r2 Microserver.....	112
Table 34:	DDR3 SODIMM Maximum Frequency Support for C32/AM3r2 Microserver.....	113
Table 35:	DDR3 UDIMM Population Support G34/C32.....	113
Table 36:	DDR3 RDIMM Population Support G34/C32.....	113
Table 37:	DDR3 LRDIMM Population Support G34/C32.....	114
Table 38:	DDR3 UDIMM Population Support AM3r2.....	114
Table 39:	DDR3 SODIMM Population Support C32/AM3r2.....	114
Table 40:	Package pin mapping.....	115
Table 41:	Phy Predriver Calibration Codes for Data/DQS at 1.5V.....	121
Table 42:	Phy Predriver Calibration Codes for Data/DQS at 1.35V.....	121

Table 43:	Phy Predriver Calibration Codes for Data/DQS at 1.25V	122
Table 44:	Phy Predriver Calibration Codes for Data/DQS at 1.5V	123
Table 45:	Phy Predriver Calibration Codes for Data/DQS at 1.35V	123
Table 46:	Phy Predriver Calibration Codes for Data/DQS at 1.25V	124
Table 47:	Phy Predriver Calibration Codes for Data/DQS at 1.5V	125
Table 48:	Phy Predriver Calibration Codes for Data/DQS at 1.35V	125
Table 49:	Phy Predriver Calibration Codes for Data/DQS at 1.25V	126
Table 50:	Phy Predriver Calibration Codes for Cmd/Addr at 1.5V	127
Table 51:	Phy Predriver Calibration Codes for Cmd/Addr at 1.35V	127
Table 52:	Phy Predriver Calibration Codes for Cmd/Addr at 1.25V	128
Table 53:	Phy Predriver Calibration Codes for Clock at 1.5V.....	129
Table 54:	Phy Predriver Calibration Codes for Clock at 1.35V.....	129
Table 55:	Phy Predriver Calibration Codes for Clock at 1.25V.....	130
Table 56:	BIOS Recommendations for RttNom and RttWr (G34r1 & UDIMM)	134
Table 57:	BIOS Recommendations for RttNom and RttWr (G34r1 & RDIMM).....	136
Table 58:	BIOS Recommendations for RttNom and RttWr (G34r1 & LRDIMM)	140
Table 59:	BIOS Recommendations for RttNom and RttWr (C32r1 & UDIMM).....	141
Table 60:	BIOS Recommendations for RttNom and RttWr (C32r1 & RDIMM).....	143
Table 61:	BIOS Recommendations for RttNom and RttWr (C32r1 & LRDIMM)	146
Table 62:	BIOS Recommendations for RttNom and RttWr (AM3r2 & UDIMM).....	146
Table 63:	BIOS Recommendations for RttNom and RttWr (SODIMM)	147
Table 64:	DDR3 UDIMM/SODIMM ODT Pattern 1 DIMM Slot/Ch	149
Table 65:	DDR3 UDIMM/SODIMM ODT Pattern 2 DIMM Slots/Ch.....	149
Table 66:	DDR3 UDIMM ODT Pattern 3 DIMM Slots/Ch	149
Table 67:	DDR3 RDIMM ODT Pattern 1 DIMM Slot/Ch	149
Table 68:	DDR3 RDIMM ODT Pattern 2 DIMM Slots/Ch.....	149
Table 69:	DDR3 RDIMM ODT Pattern 3 DIMM Slots/Ch.....	150
Table 70:	DDR3 LRDIMM ODT Pattern 1 DIMM Slot/Ch.....	150
Table 71:	DDR3 LRDIMM ODT Pattern 2 DIMM Slots/Ch	150
Table 72:	DDR3 LRDIMM ODT Pattern 3 DIMM Slots/Ch	150
Table 73:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & UDIMM)	151
Table 74:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & RDIMM)	152
Table 75:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & LRDIMM)	155
Table 76:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & UDIMM)	156
Table 77:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & RDIMM).....	158
Table 78:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & LRDIMM)	160
Table 79:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (AM3r2 & UDIMM).....	161
Table 80:	BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (SODIMM)	162
Table 81:	DCT Training Specific Register Values	163
Table 82:	BIOS Recommendations for MR0_dct[1:0][WR]	165
Table 83:	BIOS Recommendations for MR0_dct[1:0][CL[3:0]].....	165
Table 84:	BIOS Recommendations for MR2_dct[1:0][ASR, SRT].....	167
Table 85:	BIOS Recommendations for RC2[IBT] (G34r1).....	168
Table 86:	BIOS Recommendations for RC2[IBT] (C32r1	172

Table 87:	BIOS Recommendations for RC10[OperatingSpeed]	175
Table 88:	BIOS Recommendations for IBT (G34r1).....	177
Table 89:	BIOS Recommendations for IBT (C32r1).....	178
Table 90:	BIOS Recommendations for F0RC10[OperatingSpeed]	179
Table 91:	BIOS Recommendations for F0RC11[ParityCalculation].....	180
Table 92:	BIOS Recommendations for F0RC13[NumPhysicalRanks]	181
Table 93:	BIOS Recommendations for F0RC13[NumLogicalRanks].....	181
Table 94:	BIOS Recommendations for F0RC15[RankMultiplicationControl]	182
Table 95:	Register Mapping for ExtRC_x[BF:B8]	189
Table 96:	BIOS Recommendations for Write Leveling Training Seed (G34r1).....	194
Table 97:	BIOS Recommendations for Write Leveling Training Seed (C32r1).....	194
Table 98:	BIOS Recommendations for Write Leveling Training Seed (AM3r2).....	194
Table 99:	BIOS Recommendations for Receiver Enable Training Seed (G34r1)	196
Table 100:	BIOS Recommendations for Receiver Enable Training Seed (C32r1).....	197
Table 101:	BIOS Recommendations for Receiver Enable Training Seed (AM3r2).....	197
Table 102:	Command Generation and Data Comparison	204
Table 103:	BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (G34r1,UDIMM)..	205
Table 104:	BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (G34r1, RDIMM) .	206
Table 105:	BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (G34r1, LRDIMM).....	207
Table 106:	BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (C32r1, UDIMM or SODIM).....	208
Table 107:	BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (C32r1, RDIMM) .	209
Table 108:	BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (C32r1, LRDIMM)210	
Table 109:	BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (AM3r2, UDIMM or SODIM).....	211
Table 110:	Recommended interleave configurations.....	212
Table 111:	DDR3 Swapped Normalized Address Lines for CS Interleaving.....	212
Table 112:	Single Node Example.....	216
Table 113:	Dual Node Example	217
Table 114:	Dual Node with Node Interleaved Example	217
Table 115:	On-Line Spare and CS Interleaving Configurations	218
Table 116:	Link Disconnect Controls	226
Table 117:	Supported Link Operational Modes.....	228
Table 118:	MCA Register Cross-Reference Table.....	230
Table 119:	Overwrite Priorities for All Banks	233
Table 120:	Error Code Types	234
Table 121:	Error Codes: Transaction Type	235
Table 122:	Error codes: cache level	235
Table 123:	Error Codes: Memory Transaction Type.....	235
Table 124:	Error Codes: Participation Processor	235
Table 125:	Error Codes: Memory or IO.....	235
Table 126:	Error Scope Hierarchy	238
Table 127:	Recommended Scrub Rates per Node.....	240
Table 128:	Registers Commonly Used for Diagnosis.....	242

Table 129:	Configuration Specific Recommended ECC Symbol Size	245
Table 130:	x8 ECC Correctable Syndromes	246
Table 131:	x4 ECC Correctable Syndromes	251
Table 132:	Terminology in Register Descriptions	256
Table 133:	Register Mapping for D18F0x[5C:40].....	262
Table 134:	Register Mapping for D18F0x[E0,C0,A0,80]	268
Table 135:	Register Mapping for D18F0x[E4,C4,A4,84]	269
Table 136:	Register Mapping for D18F0x[E8,C8,A8,88]	272
Table 137:	Register Mapping for D18F0x[EC,CC,AC,8C].....	273
Table 138:	Register Mapping for D18F0x[F0,D0,B0,90].....	274
Table 139:	BIOS Recommendations for D18F0x[F0,D0,B0,90], D18F0x[F4,D4,B4,94]	275
Table 140:	Register Mapping for D18F0x[F4,D4,B4,94].....	276
Table 141:	Register Mapping for D18F0x[F8,D8,B8,98].....	277
Table 142:	Register Mapping for D18F0x[FC,DC,BC,9C]	277
Table 143:	Register Mapping for D18F0x[11C,118,114,110].....	278
Table 144:	Register Mapping for D18F0x[12C,128,124,120].....	278
Table 145:	Register Mapping for D18F0x[14C:130].....	279
Table 146:	Register Mapping for D18F0x[18C:170].....	284
Table 147:	Register Mapping for D18F1x[17C:140,7C:40].....	288
Table 148:	Register Mapping for D18F1x[78,70,68,60,58,50,48,40]	289
Table 149:	Register Mapping for D18F1x[178,170,168,160,158,150,148,140]	290
Table 150:	Register Mapping for D18F1x[7C,74,6C,64,5C,54,4C,44].....	290
Table 151:	Register Mapping for D18F1x[17C,174,16C,164,15C,154,14C,144].....	291
Table 152:	Register Mapping for D18F1x[1CC:180,BC:80].....	291
Table 153:	Register Mapping for D18F1x[1B8,1B0,1A8,1A0,B8,B0,A8,A0,98,90,88,80]	292
Table 154:	Register Mapping for D18F1x[1BC,1B4,1AC,1A4,BC,B4,AC,A4,9C,94,8C,84]	293
Table 155:	Register Mapping for D18F1x[1CC:1C0,19C:180].....	294
Table 156:	Register Mapping for D18F1x[D8,D0,C8,C0]	295
Table 157:	Register Mapping for D18F1x[DC,D4,CC,C4]	296
Table 158:	Register Mapping for D18F1x[EC:E0].....	296
Table 159:	Register Mapping for D18F2x[5C:40]_dct[1:0]	301
Table 160:	DIMM, Chip Select, and Register Mapping	301
Table 161:	Register Mapping for D18F2x[6C:60]_dct[1:0]	303
Table 162:	BIOS Recommendations for D18F2x[6C:60]_dct[1:0][RankDef].....	303
Table 163:	DDR3 DRAM Address Mapping.....	305
Table 164:	DDR3 LRDIMM Rank Multiplication Address Mapping.....	306
Table 165:	Valid Values for D18F2x94_dct[1:0][MemClkFreq]	310
Table 166:	Register Mapping for D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0]	315
Table 167:	Field Mapping for D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0]	315
Table 168:	Register Mapping for D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0]	317
Table 169:	Field Mapping for D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0]	318
Table 170:	Register Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]	321
Table 171:	Field Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]	322
Table 172:	Register Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]	323

Table 173:	Field Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]	323
Table 174:	Register Mapping for D18F2x9C_x0000_00[52:50]_dct[1:0]	324
Table 175:	Field Mapping for D18F2x9C_x0000_00[52:50]_dct[1:0]	324
Table 176:	Register Mapping for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0]	325
Table 177:	Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]	325
Table 178:	Index Mapping for D18F2x9C_x0D0F_0[F,8:0]04_dct[1:0]	326
Table 179:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0]	326
Table 180:	Register Mapping for D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0]	327
Table 181:	Index Mapping for D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0]	328
Table 182:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0]	328
Table 183:	Index Mapping for D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0]	328
Table 184:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0]	328
Table 185:	Index Mapping for D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0]	329
Table 186:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0]	329
Table 187:	Index Mapping for D18F2x9C_x0D0F_0[F,8:0]1F_dct[1:0]	329
Table 188:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]1F_dct[1:0]	330
Table 189:	Index Mapping for D18F2x9C_x0D0F_0[F,8:0]30_dct[1:0]	330
Table 190:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]30_dct[1:0]	330
Table 191:	Index Mapping for D18F2x9C_x0D0F_0[F,8:0]31_dct[1:0]	331
Table 192:	Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]31_dct[1:0]	331
Table 193:	Index Mapping for D18F2x9C_x0D0F_2[2:0]02_dct[1:0]	331
Table 194:	Index Mapping for D18F2x9C_x0D0F_[C,8,2][2:0]1F_dct[1:0]	332
Table 195:	Index Mapping for D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0]	333
Table 196:	Broadcast Mapping for D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0]	333
Table 197:	Index Mapping for D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0]	333
Table 198:	Index Mapping for D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0]	335
Table 199:	Index Mapping for D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0]	336
Table 200:	BIOS Recommendations for D18F2xA8_dct[1:0][CsMux45]	340
Table 201:	BIOS Recommendations for D18F2xA8_dct[1:0][CsMux67]	341
Table 202:	BIOS Recommendations for D18F2x1B0, D18F2x1B4	351
Table 203:	BIOS Recommendations for D18F2x20C_dct[1:0][Tcwl]	356
Table 204:	Register Mapping for D18F2x210_dct[1:0]_nbp[3:0]	357
Table 205:	BIOS Recommendations for D18F2x210_dct[1:0]_nbp[3:0][DataTxFifoWrDly]	357
Table 206:	BIOS Recommendations for D18F2x210_dct[1:0]_nbp[3:0][RdPtrInit]	357
Table 207:	BIOS Recommendations for D18F2x220_dct[1:0][Tmod]	361
Table 208:	BIOS Recommendations for D18F2x224_dct[1:0][Tzqcs, Tzqoper]	361
Table 209:	Register Mapping for D18F2x[234:230]_dct[1:0]	363
Table 210:	Field Mapping for D18F2x[234:230]_dct[1:0]	363
Table 211:	Register Mapping for D18F2x[23C:238]_dct[1:0]	364
Table 212:	Field Mapping for D18F2x[23C:238]_dct[1:0]	364
Table 213:	BIOS Recommendations for D18F2x248_dct[1:0][Txp] (SODIMM UDIMM)	365
Table 214:	BIOS Recommendations for D18F2x248_dct[1:0][Txp] (RDIMM LRDIMM)	366
Table 215:	BIOS Recommendations for D18F2x248_dct[1:0][Txpdll]	366
Table 216:	BIOS Recommendations for D18F2x24C_dct[1:0][Tcksr, Tcksre, Tpd]	367

Table 217:	Register Mapping for D18F2x25[8,4]_dct[1:0].....	369
Table 218:	XBAR Definitions.....	384
Table 219:	SMAF Action Definition	387
Table 220:	BIOS Recommendations for D18F3x140[FreeTok, UpRspTok].....	404
Table 221:	Register Mapping for D18F3x1[54,50,4C,48].....	405
Table 222:	BIOS Recommendations for D18F3x1[54,50,4C,48].....	406
Table 223:	Register Mapping for D18F4x1[9C,94,8C,84]_x[D0,C0].....	428
Table 224:	Register Mapping for D18F4x1[9C,94,8C,84]_x[D1,C1].....	429
Table 225:	Register Mapping for D18F4x1[9C,94,8C,84]_x[D3,C3].....	431
Table 226:	Register Mapping for D18F4x1[9C,94,8C,84]_x[D4,C4].....	432
Table 227:	BIOS Recommendations for D18F4x1[9C,94,8C,84]_x[D4,C4][DCV].....	432
Table 228:	Register Mapping for D18F4x1[9C,94,8C,84]_x[D5,C5].....	433
Table 229:	Register Mapping for D18F4x1[9C,94,8C,84]_x[DF,CF]	434
Table 230:	BIOS Recommendations for P0RcvRdPtr, P1RcvRdPtr.	434
Table 231:	BIOS Recommendations for P0XmtRdPtr, P1XmtRdPtr.....	434
Table 232:	Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]6_dm[1]	443
Table 233:	Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]6_dm[1]	444
Table 234:	Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1]	444
Table 235:	Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1].....	445
Table 236:	Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]F_dm[1].....	446
Table 237:	Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]F_dm[1]	447
Table 238:	BIOS Recommendation for DllProcessFreqCtlIndex2.....	448
Table 239:	Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][9,1]1_dm[1]	448
Table 240:	Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][9,1]1_dm[1]	449
Table 241:	Register Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]0_dm[1]	449
Table 242:	Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]0_dm[1]	450
Table 243:	Register Mapping for D18F4x1[9C,94,8C,84]_x6[9:8]84_dm[1].....	451
Table 244:	Register Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]C_dm[1]	451
Table 245:	Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]C_dm[1].....	452
Table 246:	Gen3 Deemphasis Settings.....	452
Table 247:	Estimated deemphasis (dB) for microstrip routing	452
Table 248:	Estimated deemphasis (dB) for stripline routing	453
Table 249:	D18F5x80[Enabled, DualCore] Definition.....	458
Table 250:	Register Mapping for D18F5x1[6C:60].....	461
Table 251:	NB P-state Definitions	462
Table 252:	Register Mapping for APIC[170:100]	467
Table 253:	Register Mapping for APIC[1F0:180]	468
Table 254:	Register Mapping for APIC[270:200]	468
Table 256:	Register Mapping for APIC3[60:50]	471
Table 257:	Div[3,1:0] Value Table.....	473
Table 258:	Register Mapping for APIC[4F0:480]	474
Table 259:	Register Mapping for APIC[530:500]	474
Table 260:	Reset Mapping for CPUID Fn0000_0000_E[D,C,B]X	475
Table 261:	Reset Mapping for CPUID Fn8000_0000_E[D,C,B]X	482

Table 262:	Reset Mapping for CPUID Fn8000_000[4:2]_E[D,C,B,A]X.....	485
Table 263:	ECX mapping to Cache Type for CPUID Fn8000_001D_E[D,C,B,A]X.....	495
Table 264:	Register Mapping for MSR0000_020[E,C,A,8,6,4,2,0]	506
Table 265:	Valid Values for Memory Type	506
Table 266:	Register Mapping for MSR0000_020[F,D,B,9,7,5,3,1]	507
Table 267:	Register Mapping for MSR0000_02[6F:68,59:58,50].....	508
Table 268:	Fixed-size MTRR size and Range Mapping	508
Table 269:	LS Error Descriptions	512
Table 270:	LS Error Signatures.....	513
Table 271:	LS Address Register.....	514
Table 272:	IF Error Descriptions	516
Table 273:	IF Error Signatures.....	519
Table 274:	IF Address Register.....	520
Table 275:	MBE, SBU, and SBC Definitions	524
Table 276:	CU Error Descriptions	524
Table 277:	CU Error Signatures.....	526
Table 278:	CU Address Register.....	528
Table 279:	NB Error Descriptions	533
Table 280:	NB Error Signatures, Part 1	535
Table 281:	NB Error Signatures, Part 2	536
Table 282:	NB Address Register Default Encoding	538
Table 283:	NB Address Register for Protocol Errors	538
Table 284:	Protocol Error Type.....	538
Table 285:	NB Address Register for NB Array Errors	540
Table 286:	NB Address Register for L3 Array Errors	541
Table 287:	NB Address Register for Watchdog Timer Errors	541
Table 288:	EX Error Descriptions.....	545
Table 289:	EX Error Signatures.....	546
Table 290:	EX Address Register.....	546
Table 291:	FP Error Descriptions.....	548
Table 292:	FP Error Signatures.....	548
Table 293:	Register Mapping for MSRC001_00[03:00]	555
Table 294:	Register Mapping for MSRC001_00[07:04]	555
Table 295:	Register Mapping for MSRC001_00[35:30]	561
Table 296:	BIOS Recommendations for MSRC001_00[35:30]	561
Table 297:	Register Mapping for MSRC001_00[53:50]	565
Table 298:	Register Mapping for MSRC001_00[6B:64].....	569
Table 299:	P-state Definitions.....	570
Table 300:	Register Mapping for MSRC001_020[A,8,6,4,2,0].....	580
Table 301:	Register Mapping for MSRC001_020[B,9,7,5,3,1].....	581
Table 302:	Register Mapping for MSRC001_024[6,4,2,0]	581
Table 303:	Register Mapping for MSRC001_024[7,5,3,1]	583
Table 304:	Register Mapping for PMCx0D[F:C]	614
Table 305:	Register Mapping for NBPMCx[1F9,0F8,0F7,0F6].....	620

Revision History

Revision 3.14 Changes

- Table 27: [DDR3 RDIMM Maximum Frequency Support for G34]: Updated.
- Table 28: [DDR3 LRDIMM Maximum Frequency Support for G34]: Updated.
- 2.10.5.10 [DRAM Phy Power Savings]: Updated.
- Table 25: [DCT Definitions]: Updated.
- MSRC001_001F[DisDatMsk]: Updated.

Revision 3.12 Changes

- 1.5.3 [Changes For Revision C]: Updated.
- 2.4.3 [Processor Cores and Downcoring]: Updated.
- 2.11.1 [The Tctl Temperature Scale]: Updated.
- Table 28: [DDR3 LRDIMM Maximum Frequency Support for G34]: Updated.
- Table 58: [BIOS Recommendations for RttNom and RttWr (G34r1 & LRDIMM)]: Updated.
- Table 61: [BIOS Recommendations for RttNom and RttWr (C32r1 & LRDIMM)]: Updated.
- Table 88: [BIOS Recommendations for IBT (G34r1)]: Updated.
- Table 89: [BIOS Recommendations for IBT (C32r1)]: Updated.
- Table 99: [BIOS Recommendations for Receiver Enable Training Seed (G34r1)]: Updated.
- Table 100: [BIOS Recommendations for Receiver Enable Training Seed (C32r1)]: Updated.
- Table 249 [D18F5x80[Enabled, DualCore] Definition]: Updated.
- D18F3x64[HtcPstateLimit]: Updated.
- D18F3x68[SwPstateLimit]: Updated.
- D18F3xA4[CurTmp]: Updated.
- CPUID Fn8000_001D_EAX_x1[NumSharingCache]: Updated.
- CPUID Fn8000_001D_EAX_x2[NumSharingCache]: Updated.
- MSRC001_0072[PstateLimit]: Updated.
- MSRC001_1021[DisWayFilter]: Added .
- 2.4.8.1 [Local APIC]: Updated.
- 2.10 [DRAM Controllers (DCTs)]: Updated.
- 2.11.2.1 [PROCHOT_L and Hardware Thermal Control (HTC)]: Updated.
- 2.13.2 [DRAM Considerations for ECC]: Clarified.
- Table 33, Table 34: Added.
- Table 62: [BIOS Recommendations for RttNom and RttWr (AM3r2 & UDIMM)]: Updated.
- Table 63: [BIOS Recommendations for RttNom and RttWr (SODIMM)]: Added.
- Table 64, Table 65, Table 97, Table 98, Table 100, Table 101, Table 106, Table 109, Table 162, Table 213: Updated.
- Table 79: [BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (AM3r2 & UDIMM)]: Updated.
- Table 80: [BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (SODIMM)]: Added.
- D18F1x120: Updated.
- D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0][ProcOdtAdv]: Updated.
- D18F2x250_dct[1:0][LfsrRollOver]: Added.
- D18F3x180[SyncFloodOnUsPwDataErr]: Updated.
- D18F3x180[McaLogUsPwDatErrEn]: Updated.
- MSRC000_0080[FFXSE]: Updated.
- MSRC001_024[6,4,2,0][IntCoreSel, IntCoreEn]: Added.

- MSRC001_1030[IbsFetchVal]: Updated.
- CPUID Fn0000_0001_EDX[FXSR]: Updated.
- CPUID Fn8000_0006_EDX[L3Assoc]: Updated.
- CPUID Fn8000_001D_EBX_x3[CacheNumWays]: Updated.

Revision 3.08 Changes

- D18F1x[17C:140,7C:40]: Added programming requirement.
- D18F4x15C[BoostSrc]: Updated.

Revision 3.06 Changes

- MSRC001_0077: Updated.
- 2.5.2.1.1.1 [TDP Limiting]: Updated.
- D18F0x68[ATMMModeEn]: Change recommendation of ATM to disabled for single threaded compute unit.
- D18F3x64[HtcPstateLimit]: Updated.
- Table 28 [DDR3 LRDIMM Maximum Frequency Support for G34]: Updated.
- Table 31 [DDR3 LRDIMM Maximum Frequency Support for C32]: Updated.
- Table 58 [BIOS Recommendations for RttNom and RttWr (G34r1 & LRDIMM)]: Updated.
- Table 61 [BIOS Recommendations for RttNom and RttWr (C32r1 & LRDIMM)]: Updated.
- Table 75 [BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & LRDIMM)]: Updated.
- Table 78 [BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & LRDIMM)]: Updated.
- Table 88 [BIOS Recommendations for IBT (G34r1)]: Updated.
- Table 89 [BIOS Recommendations for IBT (C32r1)]: Updated.
- F0RC0[OutputWeakDrive]: Updated.

Revision 3.04 Changes

- 2.5.2.1.6 [Core P-state Bandwidth Requirements]: Updated rule: (core COF / NB COF) <= 4.
- 2.5.2.1.9 [Processor-Systemboard Power Delivery Compatibility Check]: Updated.
- 2.10.5.3.4 [Phy Compensation Initialization], 2.10.5.8.4 [DQS Position Training]: Updated.
- MSRC001_1022[DisPfHwForSw]: Added.
- MSRC001_1029[ConvertPrefetchToNop]: Added.
- MSRC001_102B[PfcDis, PfcStrideDis, PfcRegionDis, PfcL1TrainDis]: Added.

Revision 3.02 Changes

- 1.5.1 [Revision Conventions]: Updated.

Revision 3.00 Changes

- Initial public release.

1 Overview

This document defines AMD Family 15h Models 00h-0Fh Processors, henceforth referred to as the processor.

- The processor overview is located at [2.1 \[Processor Overview\]](#).
- The processors is distinguished by the combined ExtFamily and BaseFamily fields of the CPUID instruction (see [CPUID Fn8000_0001_EAX](#) in [3.10 \[CPUID Instruction Registers\]](#)).

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.
- Software Optimization Guide for AMD Family 15h Processors, #47414.
- CPUID Specification, #25481.
- AMD Voltage Regulator Specification, #40182.
- 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).
- Advanced Platform Management Link (APML) Specification, #41918.
- AMD64 Technology Lightweight Profiling Specification, #43724.
- AMD Socket G34 Processor Functional Data Sheet, #42937.
- AMD Socket C32 Processor Functional Data Sheet, #47390.
- AMD Socket AM3 Processor Functional Data Sheet, #40778.
- AMD Family 15h Models 00h-0Fh Processor Electrical Data Sheet, #47079.
- Revision Guide for AMD Family 15h Models 00h-0Fh Processors, #48063.
- HyperTransport™ I/O Link Specification. www.hypertransport.org. Referred to as the link specification.

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 LSBs are Xlate[3:0].
	Bitwise OR operator. E.g. (01b 10b == 11b).
	Logical OR operator. E.g. (01b 10b == 1b); treats multibit operand as 1 if >=1 and produces a 1-bit result.
&	Bitwise AND operator. E.g. (01b & 10b == 00b).
&&	Logical AND operator. E.g. (01b && 10b == 1b); treats multibit operand as 1 if >=1 and produces a 1-bit result.
^	Bitwise exclusive-OR operator. E.g. (01b ^ 10b == 11b). Sometimes used as “raised to the power of” as well, as indicated by the context in which it is used. E.g. (2^2 == 4).
~	Bitwise NOT operator (also known as ones complement). E.g. (~10b == 01b).
!	Logical NOT operator. E.g. (!10b == 0b); treats multibit operand as 1 if >=1 and produces a 1-bit result.
==	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/ternary operator. E.g. condition? value-if-true: value-if-false. Equivalent to IF condition THEN value-if-true ELSE value-if-false.

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.
COUNT	COUNT(integer-expression): Returns the number of binary 1’s in the integer.

Table 2: Functions

Function	Definition
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).
POW	POW(base, exponent): POW(x,y) returns the value x to the power of y.

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].
APM	Application Power Management. See 2.5.2.1.1 [Application Power Management (APM)].
ATM	ATM: Accelerated Transition to Modified. See D18F0x68[ATMModeEn].
APML	Advanced Platform Management Link. See 2.13.4 [Sideband Interface (SBI)].
B	Byte.
BCS	Base configuration space. See 2.8 [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 2.5.1.1.1 [VID 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 [C-states].

Table 3: Definitions

Term	Definition
C1E	C1 enhanced state. Power-savings mode that is employed when all cores of a CMP processor are in the Halt state. See MSRC001_0055 [Interrupt Pending] .
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.
Channel interleaved mode	Mode in which DRAM address space is interleaved between DRAM channels. See 2.10.6 [Memory Interleaving Modes] .
Chipkill ECC	An error correcting code which can recover from DRAM device failures. See 2.13.2 [DRAM Considerations for ECC] .
CMP	Chip multi-processing. Refers to processors that include multiple cores. See 2.1 [Processor Overview] .
Coherent fabric	The coherent fabric is defined to include the memory and caches for each node and the coherent communication between all nodes. See 2.2 [System Overview] .
Coherent link or coh link	A link configured for coherent inter-processor traffic between nodes.
COF	Current operating frequency of a given clock domain. See 2.5.2 [P-states] .
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 2.5.2.1.1 [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 3.10 [CPUID Instruction Registers] .
CS	Chip select. See D18F2x[5C:40]_dct[1:0] [DRAM CS Base Address] .
DCT	DRAM controller. See 2.10 [DRAM Controllers (DCTs)] .
DCQ	DRAM controller queue.
DDR3	DDR3 memory technology. See 2.10 [DRAM Controllers (DCTs)] .
DFE	Decision feedback equalization. See D18F4x1[9C,94,8C,84]_x[D4,C4] [Link Phy DFE and DFR Control] .
DID	Divisor identifier. Specifies the post-PLL divisor used to reduce the COF. See 2.5.2 [P-states] .
Display refresh	Traffic used for display refresh in UMA systems. See 2.9.3.2.5 [Display Refresh And IFCM] .
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 2.10 [DRAM Controllers (DCTs)] .

Table 3: Definitions

Term	Definition
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 2.8 [Configuration Space].
EDS	Electrical data sheet. See 1.2 [Reference Documents].
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 2.5.2 [P-states].
GB	Gbyte or Gigabyte; 1,073,741,824 bytes.
#GP	A general-protection exception. Always produces an error code of 0. #GP(0)
GT/s	Giga-transfers per second.
HTC	Hardware thermal control. See 2.11.2.1 [PROCHOT_L and Hardware Thermal Control (HTC)].
HTC-active state	Hardware-controlled lower-power, lower-performance state used to reduce temperature. See 2.11.2.1 [PROCHOT_L and Hardware Thermal Control (HTC)].
IBS	Instruction based sampling. See 2.7 [Performance Monitoring].
IFCM	Isochronous flow-control mode, as defined in the link specification. See 2.9.3.2.5 [Display Refresh And IFCM].
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 D18F0x[18C:170] [Link Extended Control].
Inactive lane	The inactive lanes of a link are lanes between the operating width as specified by D18F0x[E4,C4,A4,84][WidthOut, WidthIn] and the maximum link width as specified by D18F0x[E4,C4,A4,84][MaxWidthOut, MaxWidthIn]. See the link specification.
IO configuration	Access to configuration space through IO ports CF8h and CFCh. See 2.8 [Configuration Space].
IO hub	The platform device that contains the bridge to the system BIOS.
IO link	A link configured for non-coherent traffic, per the link specification.
IORR	IO range register. See MSRC001_00[18,16] [IO Range Base (IORR_BASE[1:0])].
KB	Kbyte or Kilobyte; 1024 bytes.
L1 cache	The level 1 caches (instruction cache and the data cache), the level 2 caches, and level 3 cache of the core. See 2.1 [Processor Overview].
L2 cache	
L3 cache	
Linear (virtual) address	The address generated by a core after the segment is applied.
Link	Generic term that may refer to an IO link or a coherent 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].

Table 3: Definitions

Term	Definition
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 2.9 [Northbridge (NB)].
MCQ	Memory controller queue. See 2.9 [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 the software optimization guide, referenced in 1.2 [Reference Documents].
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].
MMIO configuration	Access to configuration space through memory space. See 2.8 [Configuration Space].
MSR	Model-specific register. The core includes several MSRs for general configuration and control. See 3.11 [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.
NFCM	Normal flow control mode. This is the flow control mode when neither Display Refresh nor IFCM are in use. See 2.9.3.2.5 [Display Refresh And IFCM].
Node ID	The identifier assigned to each node, D18F0x60[NodeId].
Node	See 2.1 [Processor Overview].
Normalized address	Addresses used by DCTs. See 2.9 [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].
PDS	Product data sheet.
Physical address	Addresses used by cores in transactions sent to the NB.

Table 3: Definitions

Term	Definition
PRBS	Pseudo-random bit sequence.
Probe filter	See 2.9.4.1 [Probe Filter].
Processor	See 2.1 [Processor Overview].
PSI	Power Status Indicator. See 2.5.1.4.1 [PSI_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 2.13 [RAS and Advanced Server Features].
RDQ	Read data queue.
RX	Receiver.
SBI	Sideband Interface. See 2.13.4 [Sideband Interface (SBI)].
Scrubber	Background memory checking logic. See 2.9.5 [Memory Scrubbers].
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 changing the voltage to a new value in one step (as opposed to stepping). See 2.5.1.5.1 [Hardware-Initiated Voltage Transitions].
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], 2.13.4 [Sideband Interface (SBI)], 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)].
Southbridge	Same as IO hub.
Speculative event	A performance monitor event counter that counts all occurrences of the event even if the event occurs during speculative code execution.
Sublink	An 8-bit-or-less (CAD) block of link signals of a link; each sublink of a link may connect to a different device. See 2.12 [Links].
SVI	Serial VID interface. See 2.5.1 [Processor Power Planes And Voltage Control].
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 2.11 [Thermal Functions].
Tctl	Processor temperature control value. See 2.11.2 [Temperature-Driven Logic].
TDC	Thermal design current. See the AMD Infrastructure Roadmap, #41482.

Table 3: Definitions

Term	Definition
TDP	Thermal design power. A power consumption parameter that is used in conjunction with thermal specifications to design appropriate cooling solutions for the processor. See 2.5.2.1.1.1 [TDP Limiting] .
Token	A scheduler entry used in various northbridge queues to track outstanding requests. See D18F3x140 [SRI to XCS Token Count] on Page 461.
Triple-Plane	Refers to a processor or systemboard where VDD0, VDD1, and VDDNB are separate and may operate at independent voltage levels. Refer to 2.5.1 [Processor Power Planes And Voltage Control] .
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.
Unganged	A link that has been divided into multiple sublinks.
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 2.9 [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 mnemonic of the format [OR_B2](#) instead of specific model numbers. The following table contains the definitions based on model and stepping used in this document. All behavior marked with a “>=” applies to future revisions. E.g. ([PROC](#)>=[OR_B2](#)). See the revision guide for additional information about revision determination. See [1.2 \[Reference Documents\]](#).

Table 4: Processor revision conventions

Term	Definition
PROC	Processor. PROC = {CPUID Fn0000_0001_EAX[ExtFamily], CPUID Fn0000_0001_EAX[ExtModel], CPUID Fn0000_0001_EAX[BaseModel], CPUID Fn0000_0001_EAX[Stepping]}.
OR_B2	OR_B2 = {06h,01h,2h}.
OR_C0	OR_C0 = {06h,02h,0h}.

1.5.2 Major Changes Relative to Family 10h Revision D Processors

- CPU core changes:
 - New processor core architecture based on compute units. Each compute unit implements 2 cores.
 - Integer cores share an FPU.
 - Support for up to 8 cores per node in product variations.
 - MSR Changes:
 - Some MSR's are shared between cores. See 2.4.1.1 [Registers Shared by Cores in a Compute Unit].
 - MSR0000_0000 [Load-Store MCA Address]: Changed to be an alias of MSR0000_0402 [LS Machine Check Address (MC0_ADDR)] instead of DC/MC3_ADDR (MSR0000_040E [MC3 Machine Check Address (MC3_ADDR)]) because LS/MC3 was merged with DC/MC0.
 - MSR0000_0001 [Load-Store MCA Status]: Changed to be an alias of MSR0000_0401 [LS Machine Check Status (MC0_STATUS)] instead of DC/MC3_ADDR (MSR0000_040D [MC3 Machine Check Status (MC3_STATUS)]) because LS/MC3 was merged with DC/MC0.
 - MSR0000_0010 [Time Stamp Counter (TSC)]: Deprecated TSC rate at NB P-state 0 MSRC001_0015[TscFreqSel]==0; TSC always runs at P0 frequency.
 - MSR0000_00E7 [Max Performance Frequency Clock Count (MPERF)]: Added.
 - MSR0000_00E8 [Actual Performance Frequency Clock Count (APERF)]: Added.
 - MSR0000_0179 [Global Machine Check Capabilities (MCG_CAP)]: Count changed to 07h. (7 banks)
 - MSR0000_017B [Global Machine Check Exception Reporting Control (MCG_CTL)]: MC0 repurposed to DC; DC/MC3 merged into LS/MC0; MC3 deprecated; MC5 repurposed to EX; MC6 added as FP.
 - MSR0000_020[F:0] [Variable-Size MTRRs Base/Mask]: SharedC.
 - MSR0000_02[6F:68,59:58,50] [Fixed-Size MTRRs]: SharedC.
 - MSR0000_02FF [MTRR Default Memory Type (MTRRdefType)]: SharedC.
 - MSR0000_0400 [LS Machine Check Control (MC0_CTL)]: Changed.
 - MSR0000_0401 [LS Machine Check Status (MC0_STATUS)]: Added Deferred, Poison; removed Scrub.
 - MSR0000_0403 [LS Machine Check Miscellaneous (MC0_MISC)]: Added LS thresholding.
 - MSR0000_0404 [IF Machine Check Control (MC1_CTL)]: SharedC; Changed.
 - MSR0000_0407 [IF Machine Check Miscellaneous (MC1_MISC)]: Added IC thresholding.
 - MSR0000_0408 [CU Machine Check Control (MC2_CTL)]: SharedC; changed; BU now named CU.
 - MSR0000_040B [CU Machine Check Miscellaneous (MC2_MISC)]: Added CU thresholding.
 - MSR0000_040C [MC3 Machine Check Control (MC3_CTL)]: Events moved to MC0/LS.
 - MSR0000_0410 [NB Machine Check Control (MC4_CTL)]: Changed.
 - MSR0000_0414 [EX Machine Check Control (MC5_CTL)]: Changed.
 - MSR0000_0417 [EX Machine Check Miscellaneous (MC5_MISC)]: Added EX thresholding.
 - MSR0000_0418 [FP Machine Check Control (MC6_CTL)]: SharedC; Added.
 - MSRC000_0104 [Time Stamp Counter Ratio (TscRateMsr)]: Added effective frequency interface.

- MSRC000_0105 [Lightweight Profile Configuration (LWP_CFG)], MSRC000_0106 [Lightweight Profile Control Block Address (LWP_CBADDR)]: Added LWP.
- MSRC000_040A [Machine Check Misc 4 (Thresholding) Register 3 (MC4_MISC3)]: Dropped.
- MSRC001_00[03:00] [Performance Event Select (PERF_CTL[3:0])], MSRC001_00[07:04] [Performance Event Counter (PERF_CTR[3:0])]: See 1.5.2.1 [Major Changes to Performance Counters Relative to Fam10h RevD].
- MSRC001_0010 [System Configuration (SYS_CFG)]: SharedC; Dropped ChgToDirtyDis, SetDirtyEnO, SetDirtyEnS, SetDirtyEnE.
- MSRC001_0015 [Hardware Configuration (HWCR)]:
 - Dropped TscFreqSel; TSC can no longer be selected to run at NB P0-state.
 - Dropped MisAlignSseDis; ability to disable misaligned SSE support is dropped.
 - Dropped SseDis; ability to disable SSE, SSE2, SSE3, and SSE4A is dropped.
 - Dropped LimitCpuIdStdMaxVal; ability to hide standard functions >1 is dropped.
 - Dropped PatchDis; ability to disable patches is dropped.
 - Dropped SlowFence; ability to select slow fence is dropped.
- MSRC001_00[18,16] [IO Range Base (IORR_BASE[1:0])]: SharedC.
- MSRC001_00[19,17] [IO Range Mask (IORR_MASK[1:0])]: SharedC.
- MSRC001_001A [Top Of Memory (TOP_MEM)]: SharedC.
- MSRC001_001D [Top Of Memory 2 (TOM2)]: SharedC.
- MSRC001_001F [NB Configuration 1 (NB_CFG1)]:
 - Dropped InitApicIdCpuIdLo==0; CoreNum is always in the least significant position in ApidId.
- MSRC001_00[35:30] [Processor Name String]: SharedNC.
- MSRC001_003E [Hardware Thermal Control (HTC)]: Changed to hardware P-state numbering.
- MSRC001_0045 [IF Machine Check Control Mask (MC1_CTL_MASK)]: SharedC.
- MSRC001_0046 [CU Machine Check Control Mask (MC2_CTL_MASK)]: SharedC.
- MSRC001_004A [FP Machine Check Control Mask (MC6_CTL_MASK)]: SharedC.
- MSRC001_00[53:50] [IO Trap (SMI_ON_IO_TRAP_[3:0])]: SharedNC.
- MSRC001_0054 [IO Trap Control (SMI_ON_IO_TRAP_CTL_STS)]: SharedNC.
- MSRC001_0055 [Interrupt Pending]: SharedNC. Dropped C1eOnCmpHalt; SmiOnCmpHalt.
- MSRC001_0061 [P-state Current Limit]: SharedC.
- MSRC001_0063 [P-state Status]: SharedC.
- MSRC001_00[6B:64] [P-state [7:0]]: SharedC. Added 3 P-states. Dropped NbVid and NbDid. Added NbPstate.
- MSRC001_0070 [COFVID Control]: Dropped NbVid and NbDid. Added NbPstate.
- MSRC001_0071 [COFVID Status]: SharedC, Dropped CurNbDid, Added NbPstateDis.
- MSRC001_0073 [C-state Base Address]: Added.
- MSRC001_0075 [APML TDP Limit]: Added. Note: SBI-only.
- MSRC001_0077 [Processor Power in TDP]: Added. Note: SBI-only.
- MSRC001_0078 [Power Averaging Period]: Added. Note: SBI-only.
- MSRC001_0079 [DRAM Controller Command Throttle]: Added. Note: SBI-only.
- MSRC001_020[A,8,6,4,2,0] [Performance Event Select (PERF_CTL[5:0])], MSRC001_020[B,9,7,5,3,1] [Performance Event Counter (PERF_CTR[5:0])], MSRC001_024[6,4,2,0] [Northbridge Performance Event Select (NB_PERF_CTL[3:0])], MSRC001_024[7,5,3,1] [Northbridge Performance Event Counter (NB_PERF_CTR[3:0])]: See 1.5.2.1 [Major Changes to Performance Counters Relative to Fam10h RevD].
- MSRC001_1030 [IBS Fetch Control (IC_IBS_CTL)]:
 - IbsL1TlbPgSz: Added 1 GB page.
- MSRC001_1033 [IBS Execution Control (SC_IBS_CTL)]:
 - IbsOpCurCnt, IbsOpMaxCnt: Increased size as indicated by CPUID Fn8000_001B_EAX[OpCntExt].
- MSRC001_1035 [IBS Op Data (SC_IBS_DATA)]:

- IbsRipInvalid: Added RIP invalid check as indicated by CPUID Fn8000_001B_EAX[RipInvalidChk].
- MSRC001_1037 [IBS Op Data 3 (DC_IBS_DATA, IbsOpData3)]:
 - IbsDcL2TlbHit1G: Dropped because there is no DC L2 TLB.
- MSRC001_103B [IBS Branch Target Address (BP_IBSTGT_RIP)]: Added; support indicated by CPUID Fn8000_001B_EAX[BrnTrgt].
- Instruction Set Architecture changes:
 - CPUID Fn0000_0001_ECX[AVX]: Added AVX instruction support.
 - CPUID Fn0000_0001_ECX[XSAVE, OSXSAVE], CPUID Fn0000_000D: Added XSAVE support.
 - CPUID Fn0000_0001_ECX[AES]: Added AES instruction support.
 - CPUID Fn0000_0001_ECX[SSE41, SSE42]: Added SSE4.1 and SSE4.2 instruction support.
 - CPUID Fn0000_0001_ECX[SSSE3]: Added SSSE3 instruction support.
 - CPUID Fn0000_0001_ECX[PCLMULQDQ]: Added PCLMULQDQ instruction support.
 - CPUID Fn0000_0006_ECX[EffFreq]: Added effective frequency interface.
 - CPUID Fn0000_000D_EAX_x0-CPUID Fn0000_000D_EDX_x3E: Added XSAVE/XRSTOR.
 - CPUID Fn8000_0001_EBX[BrandId]: Dropped.
 - CPUID Fn8000_0001_ECX[FMA4]: Added FMA4 instruction support.
 - CPUID Fn8000_0001_ECX[LWP], CPUID Fn8000_001C: Added lightweight profiling support.
 - CPUID Fn8000_0001_ECX[XOP]: Added XOP instruction support.
 - CPUID Fn8000_0001_EDX[3DNow, 3DNowExt]: Dropped 3DNow!TM and 3DNow!TM extensions.
 - CPUID Fn8000_0007_EDX[CPB]: Added Core Performance Boost.
 - CPUID Fn8000_0007_EDX[STC]: Deprecated.
 - CPUID Fn8000_0008_EAX[GuestPhysAddrSize]: Added.
 - CPUID Fn8000_0008_ECX[ApicIdCoreIdSize]: APIC ID size increased.
 - CPUID Fn8000_000A_EDX[TscRateMsr]: Added TSC ratio.
 - CPUID Fn8000_000A_EDX[VmcbClean]: Added VMCB clean bits.
 - CPUID Fn8000_000A_EDX[FlushByAsid]: Added flush by ASID.
 - CPUID Fn8000_000A_EDX[DecodeAssists]: Added decode assists.
 - CPUID Fn8000_001B_EAX[BrnTrgt, OpCntExt, RipInvalidChk]: Added.
 - CPUID Fn8000_001C_EAX-CPUID Fn8000_001C_EDX: Added LWP.
- L1, L2, and L3 Caches:
 - DC: 16 KB, 4-way, write-through, per-core.
 - IC: 64 KB, 2-way, shared between cores of a compute unit.
 - L2: 1 MB or 2MB (Product-specific), 16-way associative, shared between both cores of a compute unit.
 - L3: Up to 8 MB (Product-specific), up to 64-way associative (Product-specific), shared between compute units on node.
- TLB's:
 - Fam10h splintered 1 GB pages to 2 MB pages. Fam15h caches 1 GB pages without splintering.
 - D, L1TLB:
 - 4 KB: 32 entries, fully associative.
 - 2 MB: 32 entries, fully associative.
 - 1 GB: 32 entries, fully associative.
 - D, L2TLB:
 - None. Full size of unified TLB reported as L2 DTLB.
 - I, L1TLB:
 - 4 KB: 48 entries, fully associative.
 - 2 MB, 1 GB: 24 entries, fully associative. 2M and 1G entries share the same L1TLB bank.
 - I, L2TLB:
 - 4 KB: 512 entries, 4-way associative. (Same as Fam10h)
 - 2 MB: None. (Same as Fam10h)
 - 1 GB: None.

- Unified TLB:
 - 1024 entries, 8-way associative, any entry can cache:
 - D: 4K, 2M, 4M, or 1G translation.
 - I: 2M, 4M, or 1G translation. Not 4K ITLB translations.
 - Notes: Unified TLB natively stores 4M translations. An entry allocated by one core is not visible to the other core of a compute unit.
- Memory controller (MCT) and DRAM controllers (DCTs) additions:
 - Support for 933 MHz (1866 MT/s) MEMCLK frequency.
 - Low-voltage DDR3 support; Support to run DDR3 at 1.35 V and 1.25 V for power savings.
 - LRDIMM support; 1-3 LRDIMMs supported per channel.
 - DRAM Prefetcher:
 - Increased DRAM prefetch table size from 16 to 32.
 - Differentiate between core prefetch requests and core demand requests.
 - DCT register access changed. See [2.10.1](#).
 - Hardware generated DRAM training patterns. See [2.10.5.8.6.1](#).
 - Dropped support for DRAM channel ganging.
 - Dropped support for DDR2.
 - [D18F2x110\[DctSelIntLvAddr\]==01b](#) not supported if node interleaving is enabled.
- Links and IO additions:
 - Upstream EOI handled as a protocol violation.
 - Link mapping changed in G34r1 package. See [2.12.1.5](#).
- RAS-related additions:
 - Added MCA bank for FP errors. See [2.13.1](#).
 - Data poisoning and deferred errors.
 - Core error thresholding via polling.
 - [MCI_STATUS\[Overflow\]](#) no longer indicates a fatal condition. See [MSR0000_0401\[Overflow\]](#).
- General Northbridge changes:
 - Added Accelerated Transition to Modified protocol in L3, uses the Muw (modified-unwritten) state.
 - Probe filter associativity increased from 4 to 8 way.
 - Removed extended MMIO map (F1x114_x[3:2]). Added 4 base/mask pairs to the base MMIO map and changed all base MMIO map entries to support PhysAddr[47:40]. See [D18F1x\[1CC:180,BC:80\]](#).
 - Victim distribution to improve 2P G34 routing. See [2.9.3.2.4.3 \[Victim Distribution Mode\]](#).
 - GART cache size reduced from 32 entries to 16 entries. See [2.9.2 \[GART\]](#).
- Power management:
 - Support for up to 8 P-states.
 - Support for up to 2 NB P-states in a multi-node system.
 - **APM**: Allows the processor to deterministically provide maximum performance while remaining within the specified power delivery and removal envelope. See [2.5.2.1.1 \[Application Power Management \(APM\)\]](#).
 - IO based C-state interface; 3 IO based C-states supported. See [2.5.3.1](#).
 - Core C6 (CC6) support. See [2.5.3.3.3 \[Core C6 \(CC6\) State\]](#).
 - C-state Boost. See [D18F4x16C \[APM TDP Control\]](#).
 - Added IDLE_EXIT_L support for MT C1E for the AM3r2 package.
 - Effective frequency interface. See [CPUID Fn0000_0006_ECX\[EffFreq\]](#).
 - DRAM Power Management:
 - The Command/Address/Bank buses tri-state when deselected.
 - Added software throttle control over DRAM bandwidth. See [2.10.10](#).
 - Dropped PVI regulator interface.
- Core/NB performance counters:
 - See [1.5.2.1 \[Major Changes to Performance Counters Relative to Fam10h RevD\]](#).

1.5.2.1 Major Changes to Performance Counters Relative to Fam10h RevD

- Core performance counters:
 - Added 2 core performance counters; added new MSR location at [MSRC001_020\[A,8,6,4,2,0\]](#); alias [MSRC001_00\[03:00\]](#) to [MSRC001_020\[6,4,2,0\]](#). See [MSRC001_020\[A,8,6,4,2,0\]](#).
 - Added multi-event support. See 2.7.1 [Core Performance Monitor Counters].
 - Not all performance monitor events can be counted on all counters. See 2.7.1 [Core Performance Monitor Counters].
 - Changed [PMCx000 \[FPU Pipe Assignment\]](#).
 - Changed [PMCx001 \[FP Scheduler Empty\]](#).
 - Dropped [PMCx002 \[Dispatched Fast Flag FPU Operations\]](#).
 - Changed [PMCx003 \[Retired Floating Point Ops\]](#).
 - Changed [PMCx004 \[Number of Move Elimination and Scalar Op Optimization\]](#).
 - Changed [PMCx005 \[Retired Serializing Ops\]](#).
 - Changed [PMCx006 \[Number of Cycles that a Bottom-Execute uop is in the FP Scheduler\]](#).
 - Changed [PMCx023 \[Load Queue/Store Queue Full\]](#).
 - Changed [PMCx02A \[Canceled Store to Load Forward Operations\]](#).
 - Added [PMCx030 \[Executed CLFLUSH Instructions\]](#).
 - Changed [PMCx041 \[Data Cache Misses\]](#).
 - Changed [PMCx042 \[Data Cache Refills from L2 or System\]](#).
 - Changed [PMCx043 \[Data Cache Refills from System\]](#).
 - Dropped [PMCx044 \[Data Cache Lines Evicted\]](#).
 - Changed [PMCx045 \[Unified TLB Hit\]](#).
 - Changed [PMCx046 \[Unified TLB Miss\]](#).
 - Dropped [PMCx048 \[Microarchitectural Late Cancel of an Access\]](#).
 - Dropped [PMCx049 \[Microarchitectural Early Cancel of an Access\]](#).
 - Dropped [PMCx04A \[Single-bit ECC Errors Recorded by Scrubber\]](#).
 - Dropped [PMCx04C \[DCACHE Misses by Locked Instructions\]](#).
 - Dropped [PMCx04D \[L1 DTLB Hit\]](#). Covered by [PMCx045 \[Unified TLB Hit\]](#).
 - Dropped [PMCx052 \[Ineffective Software Prefetches\]](#): Software prefetch hit in L2.
 - Changed [PMCx067 \[Data Prefetcher\]](#).
 - Added [PMCx067 \[Data Prefetcher\]\[5\]](#): Modified unwritten.
 - Changed [PMCx07D \[Requests to L2 Cache\]\[6:5\]](#): Added L2 cache prefetcher request; dropped DC hardware prefetch request.
 - Changed [PMCx07E \[L2 Cache Misses\]\[4:3\]](#): Added L2 cache prefetcher request; dropped DC hardware prefetch request.
 - Added [PMCx07E \[L2 Cache Misses\]\[2\]](#): Added L2 Clean Writebacks to system.
 - Added [PMCx16C \[L2 Prefetcher Trigger Events\]](#).
 - Added [PMCx085 \[L1 ITLB Miss, L2 ITLB Miss\]\[2\]](#): Instruction fetches to a 1 GB page.
 - Added [PMCx08C \[Instruction Cache Lines Invalidated\]\[3:2\]](#): SMC that hit/missed on in-flight instructions.
 - Dropped [PMCx0CC \[Retired Fastpath Double Op Instructions\]](#).
 - Dropped [PMCx0D2 \[Dispatch Stall for Branch Abort to Retire\]](#).
 - Dropped [PMCx0D4 \[Dispatch Stall for Segment Load\]](#).
 - Dropped [PMCx0DA \[Dispatch Stall for Far Transfer or Resync to Retire\]](#).
 - Changed [PMCx0DB \[FPU Exceptions\]](#).
 - Dropped [PMCx1C0 \[Retired x87 Floating Point Operations\]](#).
 - Changed [PMCx1CF \[Tagged IBS Ops\]](#).
 - Dropped [PMCx1D3 \[LFENCE Instructions Retired\]](#).
 - Dropped [PMCx1D4 \[SFENCE Instructions Retired\]](#).
 - Dropped [PMCx1D5 \[MFENCE Instructions Retired\]](#).

- Added PMCx1D8 [Dispatch Stall for STQ Full].
- NB performance counters:
 - Added 4 NB performance counters. See MSRC001_024[6,4,2,0].
 - Moved NB performance events to NB performance counters.
 - Dropped NBPMCx0E8 [Thermal Status][4:3]: STC deprecated.
 - Changed NBPMCx0EE [GART Events]: Dropped DEV.
 - Changed NBPMCx4E0 [Read Request to L3 Cache]-NBPMCx4ED [Non-canceled L3 Read Requests]: Logical compute unit. Fam10h is physical compute unit
 - Added NBPMCx4EF [L3 Latency].

1.5.3 Changes For Revision C

See the following references for information about changes for this revision.

- Changes that may result in BIOS modifications.
 - Instruction Set Architecture changes:
 - Added TBM, BMI1, FMA, and F16C instruction support. See CPUID Fn8000_0001_ECX[TBM], CPUID Fn0000_0007_EBX_x0[BMI1], CPUID Fn0000_0001_ECX[FMA, F16C].
 - Added selective PDC invalidation on INVLPG; added MSRC000_0080[TCE] and CPUID Fn8000_0001_ECX[TCE].
 - Addition of read only APERF/MPERF MSRs to allow for sharing. Added CPUID Fn8000_0007_EDX[EffFreqRO], MSRC000_00E7, MSRC000_00E8, MSRC001_0015[EffFreqReadOnlyLock].
 - Increased L1 DTLB size from 32 to 64 entries. Affects CPUID Fn8000_0005_EBX[L1DTIb4KSize], CPUID Fn8000_0005_EAX[L1DTIb2and4MSize], and CPUID Fn8000_0019_EAX[L1DTIb1GSize].
 - 2.3.3 [Using L2 Cache as General Storage During Boot]: Updated.
 - 2.3.3.1 [Using L2 Cache as General Storage During Boot for (PROC<OR_C0)]: Added.
 - 2.4.3 [Processor Cores and Downcoring]: selected support for downcoring a single core of a compute unit including software downcoring.
 - D18F0x150[HtRetryCrcDatIns]: Updated.
 - MSRC001_1021[DisLoopPredictor]: Added. Added BIOS recommendation to disable.
 - MSRC001_1028[DiDtCfg3], D18F3x1FC[DiDtCfg3]: Deprecated.
 - MSRC001_1028[DiDtCfg5, DiDtCfg4], D18F3x1FC[DiDtCfg5, DiDtCfg4]: Added.
 - MSRC001_102B[PwcDisableWalkerSharing]: Changed BIOS recommendation.
 - MSRC001_102B[PfcDoubleStride]: Added BIOS recommendation.
 - MSRC001_102B[PfcStrideMul]: Added.
 - MSRC001_102C[LateSbzResync]: Added.
 - MSRC001_102D [Load-Store Configuration 2 (LS_CFG2)]: Added.

1.5.3.1 Performance monitor changes for revision C

- Performance monitor changes:
 - PMCx032 [Misaligned Stores]: Added.
 - PMCx034 [FP +Load Buffer Stall]: Added.
 - PMCx0D[F:C] [DR[3:0] Breakpoint Matches]: Data breakpoint counting changed.
 - PMCx1C0 [Retired x87 Floating Point Operations]: Added.

2 Functional Description

2.1 Processor Overview

The *processor* is defined as follows:

- The processor is a package that contains one or more nodes.
- Supports x86-based instruction sets.
- A *node*, is an integrated circuit device that includes:
 - One to 8 cores (one to four compute units).
 - Includes up to four high-speed communication interfaces (referred to as *links*) that may be configured for HyperTransport™ technology (referred to as *IO links*) or for AMD-proprietary inter-processor communication (referred to as *coherent links*)
 - Up to two 64-bit DDR3 interfaces for communication to system memory.
 - One communication packet routing block referred to as the *northbridge* (NB).

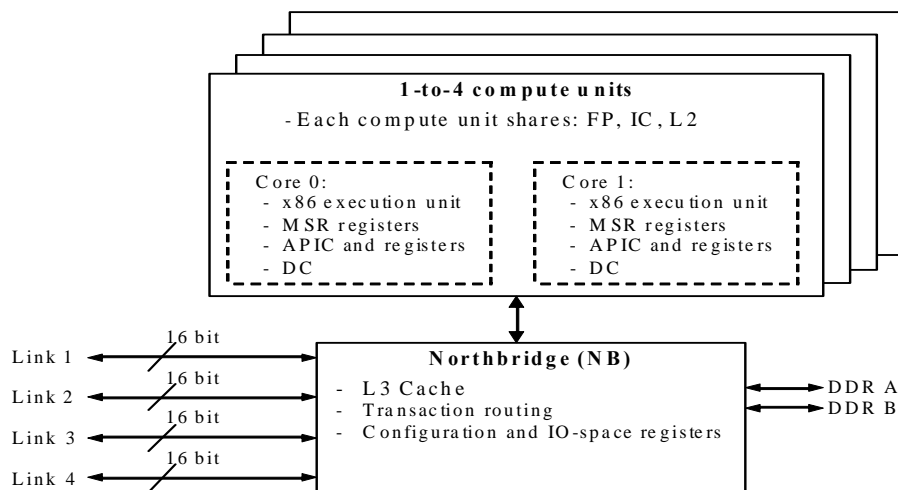


Figure 1: A Single-Node Processor

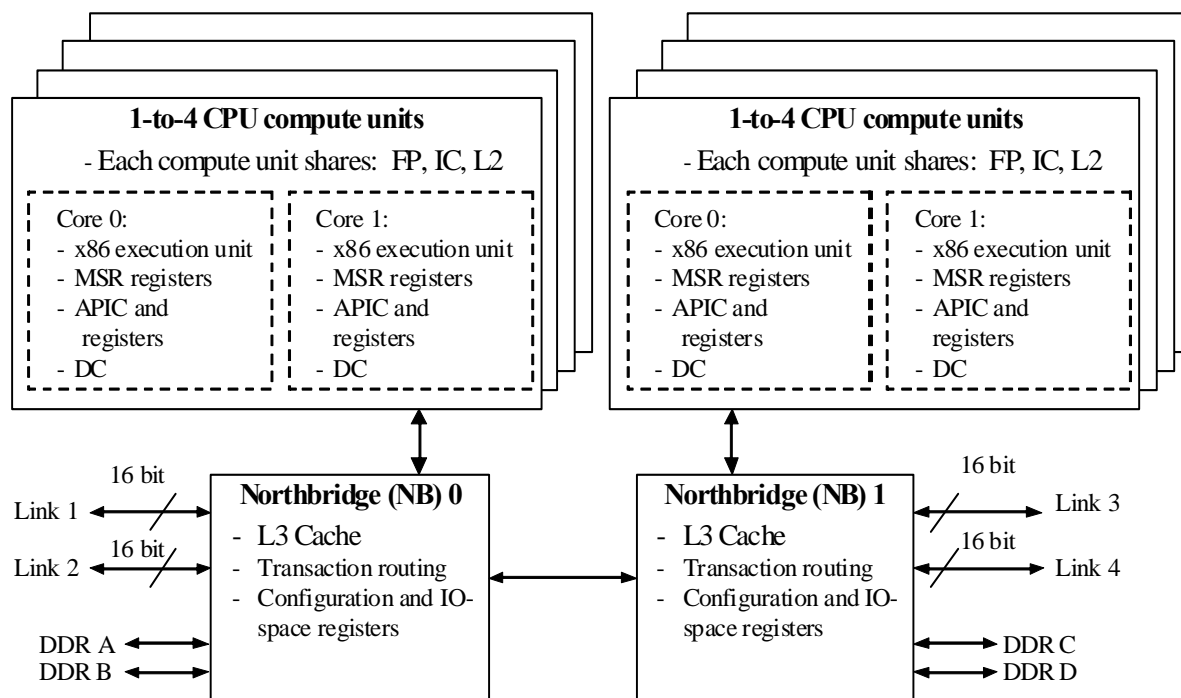


Figure 2: A Dual-Node Processor

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, first-level instruction cache, are shared between both cores of the 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.

Each link can be configured to operate under the rules of one of the following interface specifications: (1) AMD proprietary, coherent inter-processor link; or (2) non-coherent HyperTransport™ IO link. When a link is configured for non-coherent IO traffic, it is referred to as an IO link.

Each DRAM interface supports a 64-bit memory channel.

The NB routes transactions between the cores, the link, and the DRAM interfaces. It includes the configuration register space for the device. It may include an L3 cache as well.

2.2 System Overview

The following diagram illustrates the expected system architecture. Smaller systems may not include multiple processors or multiple IO links. Larger systems may include many more processors. Each processor in the coherent fabric communicates with other processors through the coherent link protocol. Processors communicate with the IO subsystem through IO links.

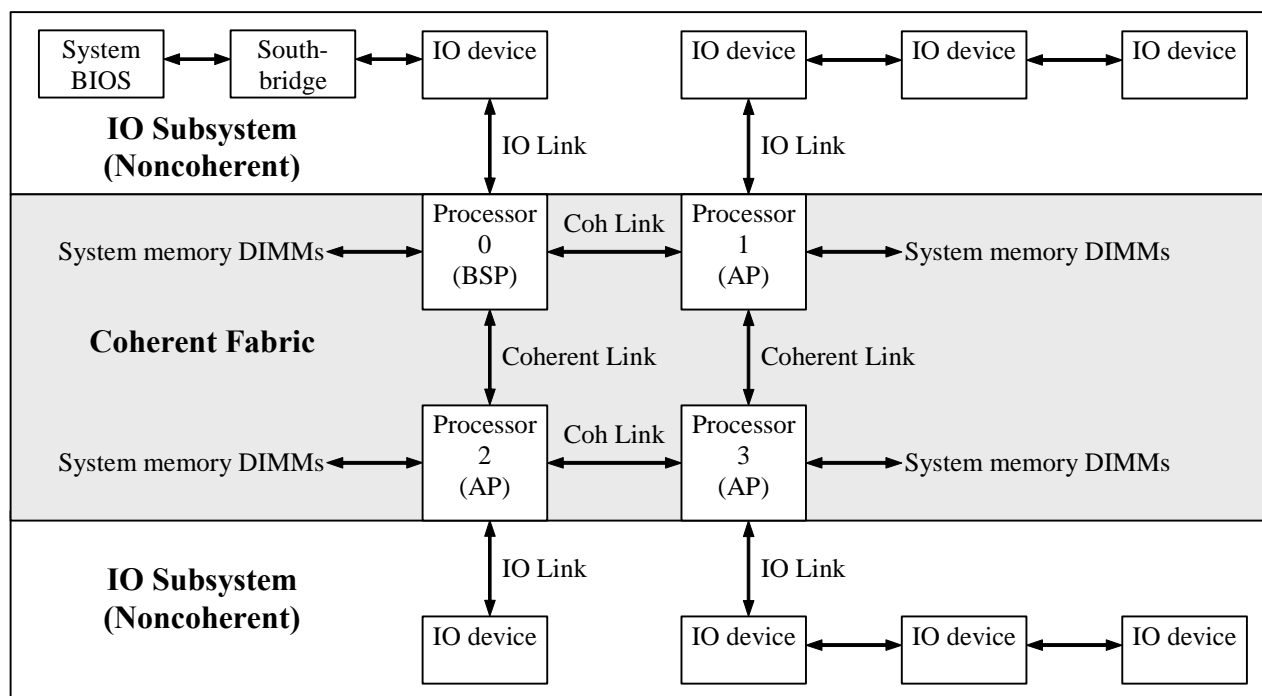


Figure 3: System Diagram

2.3 Processor Initialization

This section describes the initialization sequence after a cold reset.

The processor that is connected to the IO hub is the BSP. Core 0 of the BSP, called the BSC, begins executing code from the reset vector. Core 0 on all other nodes do not fetch code until `D18F0x6C[ReqDis]` for that node is cleared. The remaining cores do not fetch code until their enable bits are set (`D18F0x1DC[CpuEn]`).

2.3.1 BSP initialization

The BSP 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 cold reset then BIOS must clear `MSR0000_0411 [NB Machine Check Status (MC4_STATUS)]`.
 - If this is a warm reset then BIOS may check for valid MCA errors and if present save the status for later use. See 2.13.1.6 [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 of APIC (2.4.8.1.3 [ApicId Enumeration Requirements]).
- Perform coherent link enumeration (routing table and NodeID), as described in 2.9.3.2 [HyperTransport™ Technology Routing].
- Configure all IO link devices.
 - Set configuration-base and -limit (`D18F1x[EC:E0] [Configuration Map][BusNumBase, BusNum-Limit]`).

- Perform device enumeration for all IO link devices (see link specification).
- If required, reallocate data and flow control buffers of the links (see [D18F0x\[F0,D0,B0,90\]](#) [Link Base Channel Buffer Count] and [D18F0x\[F4,D4,B4,94\]](#) [Link Isochronous Channel Buffer Count]).
- Configure link speed and width. See link specification.
- Issue system warm reset.
- Configure processor power management. See [2.4 \[Processor Core\]](#).
- If supported, allow other cores to begin fetching instructions by clearing [D18F0x6C\[ReqDis\]](#) in the PCI configuration space of all other nodes and setting [D18F0x1DC\[CpuEn\]](#) in the PCI configuration space of all nodes. See [2.4.3 \[Processor Cores and Downcoring\]](#).

2.3.2 AP initialization

All other processor cores other than core 0 of node 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 type of startup using [D18F0x6C\[ColdRstDet\]](#):
 - If this is a cold reset then BIOS must clear [MSR0000_0411 \[NB Machine Check Status \(MC4_STATUS\)\]](#).
 - If this is a warm reset then BIOS may check for valid MCA errors and if present save the status for later use. See [2.13.1.6 \[Handling Machine Check Exceptions\]](#).
- Set up the local APIC. See [2.4.8.1.3 \[ApicId Enumeration Requirements\]](#).
- Configure processor power management. See [2.4 \[Processor Core\]](#).

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. [CPUID Fn8000_0006_ECX\[L2Assoc\]=8h](#) (16 ways) 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[47:16]`, `L2WayIndex=PhysAddr[15:6]`.
 - `L2Size=2 MB`: `L2Tag=PhysAddr[47:17]`, `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 [2.10.5.8.6.1 \[DRAM Training Pattern Generation\]](#).

If (PROC<OR_C0) then BIOS must also follow the requirements in [2.3.3.1 \[Using L2 Cache as General Storage During Boot for \(PROC<OR_C0\)\]](#).

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.
- If ([MSRC001_102B](#)[CombineCr0Cd]==1) Then [MSRC001_102B](#)[CombineCr0Cd]=0. See [MSRC001_102B](#)[CombineCr0Cd].
- 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 [CAR](#) memory region, marked as WB-DRAM, must be marked as UC memory type. This prevents speculative accesses that fall outside of the [CAR](#) memory region from getting to the caches and DRAM controller.
- If ([MSRC001_1023](#)[L2WayLock]==1) then:
 - Only the ways 0-([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 ((PROC>=OR_C0) && (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 ((PROC>=OR_C0) && (MSRC001_1023[L2WayLock]==1) && (MSRC001_1023[L2FirstLockedWay]==Fh)) then program MSRC001_1023[L2WayLock]=0.

2.3.3.1 Using L2 Cache as General Storage During Boot for (PROC<OR_C0)

If (PROC<OR_C0) then BIOS must also follow the requirements in this section in addition to the requirements stated in section 2.3.3 [Using L2 Cache as General Storage During Boot].

BIOS must ensure that no WB-DRAM region for one L2 on a node overlaps with a WB-DRAM region for another L2 on the same node.

2.3.4 Multiprocessing Capability Detection

The multiprocessing capability of the processor is determined by D18F3xE8[MPCap].

During boot, the BIOS checks the multiprocessing capability of all processors, and configures the system accordingly.

Multiprocessing capability detection is not required in a single processor system.

All processors must be dual-processor (DP) capable or multiprocessor (MP) capable in a DP system. If any processor is not at least DP capable, the BIOS must configure the BSP as a uni-processor (UP), and must not initialize the AP.

All processors must be MP capable in an MP system. If any processor is not MP capable, the BIOS must configure the BSP as a UP processor, and must not initialize APs.

If all processors do not have adequate multiprocessing capability for a DP or an MP system, the BIOS must display the following message:

```
***** Warning: non-MP Processor *****
```

The processor(s) installed in your system are not multiprocessing capable. Now your system will halt.

If all processors have adequate multiprocessing capability for a DP or an MP system, but have different model numbers or operate at different frequencies, see [2.5.2.3 \[Mixed Frequency and Power P-state Configuration\]](#).

2.3.5 SLIT and SRAT

The System Locality Distance Information Table (SLIT) and System Resource Affinity Table (SRAT) are described in the *Advanced Configuration and Power Interface Specification*.

2.3.5.1 SLIT

The SLIT table is programmed with the following requirements:

- The local node in the SLIT table is 10.
- For fully connected system topologies, the remaining table entries are 16. A fully connected system topology is one where the number of hops between any two nodes in the system is one.
- For system topologies that are not fully connected:
 - If probe filter is enabled (**PrbFltrEn**), remaining table entries are programmed with the formula $10 + \text{num_hops} * 6$.
 - If probe filter is disabled (\sim **PrbFltrEn**), remaining table entries are programmed such that the maximum hop entries have a value of 13 and all other entries have 10.

Consider the 8 node system topology in [Figure 4](#). The maximum number of hops between any two nodes is 3. [Table 5](#) and [Table 6](#) show the SLIT table entries without probe filter and with probe filter. The values are normalized latencies and not absolute latencies. The raw latency of a system without the probe filter enabled is higher than one with the probe filter enabled.




Figure 4: Example 8-Node System in Twisted Ladder Topology

Table 5: SLIT Table Without Probe Filter

Node	0	1	2	3	4	5	6	7
0	10	10	10	10	10	13	10	13
1	10	10	10	10	10	10	10	10
2	10	10	10	10	10	13	10	13
3	10	10	10	10	10	10	10	10
4	10	10	10	10	10	10	10	10
5	13	10	13	10	10	10	10	10
6	10	10	10	10	10	10	10	10
7	13	10	13	10	10	10	10	10

Table 6: SLIT Table With Probe Filter

Node	0	1	2	3	4	5	6	7
0	10	16	16	22	22	28	22	28
1	16	10	22	22	16	22	16	22
2	16	22	10	16	22	28	22	28
3	22	22	16	10	16	22	16	22
4	22	16	22	16	10	16	22	22
5	28	22	28	22	16	10	22	16
6	22	16	22	16	22	22	10	16
7	28	22	28	22	22	16	16	10

2.3.5.2 SRAT

A unique proximity domain is assigned for each node in the system. A processor local APIC affinity structure is created such that all cores in a node are assigned the same proximity domain as the node. The base address and length of the memory attached to the node is programmed into a memory affinity structure. See the *Advanced Configuration and Power Interface Specification* for additional information. The SRAT does not need to be created if node interleaving is enabled. See 2.10.6.3 [Node Interleaving].

2.4 Processor 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.

Table 7: Compute Unit Definitions

Term	Definition
NumOfCompUnitsOnNode	The number of compute units on a node for which at least 1 core is enabled. NumOfCompUnitsOnNode = COUNT(D18F5x80[Enabled]).
DualCoreEnabled	Both cores of a compute unit are enabled. DualCoreEnabled = (D18F5x80[DualCore[0]]==1). 0=Core 0 enabled, Core 1 disabled.

2.4.1.1 Registers Shared by Cores in a Compute Unit

Some MSRs are implemented one copy per compute unit instead of per core; these MSRs 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** MSRs 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.

The **SharedC** and **SharedNC** MSR's are listed as follows:

Table 8: SharedC and SharedNC MSR's

Register	Shared
MSR0000_020[F:0] [Variable-Size MTRRs Base/Mask]	SharedC
MSR0000_02[6F:68,59:58,50] [Fixed-Size MTRRs]	SharedC
MSR0000_02FF [MTRR Default Memory Type (MTRRdefType)]	SharedC
MSR0000_0404 [IF Machine Check Control (MC1_CTL)]	SharedC
MSR0000_0408 [CU Machine Check Control (MC2_CTL)]	SharedC
MSR0000_0418 [FP Machine Check Control (MC6_CTL)]	SharedC
MSRC001_0010 [System Configuration (SYS_CFG)], except MtrrFixDramModEn.	SharedC
MSRC001_00[18,16] [IO Range Base (IORR_BASE[1:0])]	SharedC
MSRC001_00[19,17] [IO Range Mask (IORR_MASK[1:0])]	SharedC
MSRC001_001A [Top Of Memory (TOP_MEM)]	SharedC
MSRC001_001D [Top Of Memory 2 (TOM2)]	SharedC
MSRC001_00[35:30] [Processor Name String]	SharedNC
MSRC001_0045 [IF Machine Check Control Mask (MC1_CTL_MASK)]	SharedC
MSRC001_0046 [CU Machine Check Control Mask (MC2_CTL_MASK)]	SharedC
MSRC001_004A [FP Machine Check Control Mask (MC6_CTL_MASK)]	SharedC
MSRC001_00[53:50] [IO Trap (SMI_ON_IO_TRAP_[3:0])]	SharedNC
MSRC001_0054 [IO Trap Control (SMI_ON_IO_TRAP_CTL_STS)]	SharedNC
MSRC001_0055 [Interrupt Pending]	SharedNC
MSRC001_00[6B:64] [P-state [7:0]]	SharedC
MSRC001_1021 [Instruction Cache Configuration (IC_CFG)]	SharedC
MSRC001_1023 [Combined Unit Configuration (CU_CFG)]	SharedC
MSRC001_1028 [Floating Point Configuration (FP_CFG)]	SharedC
MSRC001_102A [Combined Unit Configuration 2 (CU_CFG2)]	SharedC
MSRC001_102B [Combined Unit Configuration 3 (CU_CFG3)]	SharedC

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

Each node supports downcoring as follows:

- The number of cores supported by a node is specified by [D18F5x84\[CmpCap\]](#).
- The cores of a compute unit are even core x and odd core $x+1$.
- The cores of a compute unit may be software downcored by [D18F3x190\[DisCore\]](#) if ([DualCoreEnabled](#)==1). If ([DualCoreEnabled](#)==0) then the cores of a compute unit may not be software downcored. 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.
 - Exception:
 - The odd core of a compute-unit can be [software] downcored without downcoring the even core for 16-core processors in the G34 package as long as the silicon revision is not [OR_B2](#).
 - 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 on each node.
 - For dual-node processors, the number of cores enabled (not downcored) in each internal node must be the same. See [D18F3xE8\[IntNodeNum\]](#) for how to identify the internal nodes in a dual-node processor. The specific cores enabled in each internal node do not need to be the same, only the total number of cores in each internal node.
 - BIOS should configure all processors in a system to have the same number of enabled cores.
 - [D18F3x190\[DisCore\]](#) affects the value of [CPUID Fn8000_0008_ECX\[NC\]](#), [D18F5x80\[Enabled, Dual-Core\]](#). [D18F3x190\[DisCore\]](#) does not affect the value of [D18F5x84\[CmpCap\]](#).
- 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.
- In dual-node processors, the *CpuCoreNum* assignments are described as follows. [CPUID Fn8000_0008_ECX\[NC\]](#)+1 is the number of cores in the processor. [D18F3xE8\[IntNodeNum\]](#) specifies the internal node number. Cores 0 through $(NC+1)/2-1$ are assigned to internal node 0 and cores $(NC+1)/2$ through NC are assigned to internal node 1.

Some legacy operating systems do not support processors with a non-power-of-2 number of cores. The BIOS should support a user configurable option to disable cores down to a power-of-2 number of cores for legacy operating system support.

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

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

[D18F3x190\[DisCore\]](#) programming rules:

- Setting bits corresponding to cores that are not present results in undefined behavior.
- 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.
- Software may remove cores only once. If software removes cores by setting `D18F3x190[DisCore]=1`, then software is not allowed to disable additional cores after the next warm reset.
- The most significant bit N is (the number of cores on the node)-1 at cold reset; the number of cores on the node at cold reset is $((\text{CPUID Fn8000_0008_ECX}[NC]+1)/(\text{the number of nodes on the processor}))$; See `D18F3xE8[MultiNodeCpu]`.
- 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.
 - 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.
- If `D18F3x190[DisCore]` is changed, then the following need to be updated:
 - `D18F0x60[CpuCnt[4:0]]`, `D18F0x160[CpuCnt[7:5]]`.
 - `D18F3x1A0[L3FreeListCBC]`. Note that L3FreeListCBC is Reset-applied.
 - `D18F5x170[NbPstateThreshold]`
 - `MSRC001_102A[ThrottleNbInterface]`

2.4.4 Physical Address Space

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

- IO link requests with non-zero `PhysAddr[63:48]`.
- IO link or CPU requests with non-zero `PhysAddr[47:40]` where `D18F0x68[CHtExtAddrEn]=0`.
- IO link or CPU requests with non-zero `PhysAddr[47:40]` which targets an IO link for which the appropriate `D18F0x[E4,C4,A4,84][Addr64BitEn]=0`.
- IO link requests with non-zero `PhysAddr[47:40]` received from an IO link for which the appropriate `D18F0x[E4,C4,A4,84][Addr64BitEn]=0`.

2.4.5 System Address Map

System software must not map memory in the reserved HyperTransport™ address regions. The link specification details the address map available to system hosts and devices. Downstream host accesses to reserved HyperTransport™ address regions result in a page fault, whether in real mode or protected mode. Upstream system device accesses to reserved HyperTransport™ address regions result in 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 an IO link are routed through the NB. An IO link access to physical address space indicates to the NB the cache attribute (Coherent or Non-coherent, based on bit[0] of the Sized Read and Write commands).

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 [2.8 \[Configuration Space\]](#).

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 MTRRs. (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 [2.8 \[Configuration Space\]](#).
6. NB address space routing. See [2.9.3.1.1 \[DRAM and MMIO Memory Space\]](#).

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 [2.5.2.1.2.1 \[Software P-state Numbering\]](#) and [MSRC001_00\[6B:64\] \[P-state \[7:0\]\]](#).
- The APIC timer ([APIC380](#) and [APIC390](#)), which increments at the rate of CLKIN; 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.
- MSRC001_102B [Combined Unit Configuration 3 (CU_CFG3)] 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.

Performance losses may occur if interrupts are generated at a rate faster than which they can be serviced.

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 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 and processors in the system 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]`.) $MNC = (2 \wedge \text{CPUID Fn8000_0008_ECX[ApicIdCoreIdSize]})$. 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* on processor *i* must be enumerated/assigned as:

- $\text{ApicId[proc=i, core=j]} = (\text{OFFSET_IDX} + i) * \text{MNC} + j$
- 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 $N = (\text{Number_Of_Processors} * \text{MNC})$ and $M = \text{Number_Of_IOAPICs}$:

- If $(N+M) < 16$, then assign the local (core) ApicId's first from 0 to N-1, and the IOAPIC IDs from N to N+(M-1). APIC ID 15 is reserved for broadcast when `APIC410[ExtApicIdEn]==0`.
- If $(N+M) \geq 16$, then assign the IOAPIC IDs first from 0 to M-1, and the local (core) ApicId's from K to K+(N-1), where K is an integer multiple of MNC greater than M-1.

For example, consider a 3 processor system where each processor has 3 cores and there are 8 IOAPIC devices. Each core can support an 8-bit ApicId. But if each IOAPIC device supports only a 4-bit IOAPIC ID, then the problem can be solved by shifting the core ApicId space to start at some integer multiple of MNC, such as offset 8 ($MNC = 4$; $\text{OFFSET_IDX}=2$):

$\text{ApicId[proc=0,core=0]} = (2+0)*4 + 0 = 0x08$	$\text{ApicId[proc=1,core=2]} = (2+1)*4 + 2 = 0x0E$
$\text{ApicId[proc=0,core=1]} = (2+0)*4 + 1 = 0x09$	$\text{ApicId[proc=2,core=0]} = (2+2)*4 + 0 = 0x10$
$\text{ApicId[proc=0,core=2]} = (2+0)*4 + 2 = 0x0A$	$\text{ApicId[proc=2,core=1]} = (2+2)*4 + 1 = 0x11$
$\text{ApicId[proc=1,core=0]} = (2+1)*4 + 0 = 0x0C$	$\text{ApicId[proc=2,core=2]} = (2+2)*4 + 2 = 0x12$
$\text{ApicId[proc=1,core=1]} = (2+1)*4 + 1 = 0x0D$	

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 in `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])].
- In the NB, as specified by:
 - [D18F3xB0](#) [On-Line Spare Control].
 - [D18F3x160](#) [NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)].
- 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 9: 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
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 10: 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	

Table 10: SMM Save State

Offset	Size	Contents		Access
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	
FE80h	4 Bytes	IDTR	Reserved	Read-only
FE84h	Word		Limit	
FEB6h	2 Bytes		Reserved	
FE88h	Quadword		Base	
FE90h	Word	TR	Selector	Read-only
FE92h	Word		Attributes	
FE94h	Doubleword		Limit	
FE98h	Quadword		Base	
FEA0h	Quadword	IO_RESTART_RIP		Read-only
FEA8h	Quadword	IO_RESTART_RCX		
FEB0h	Quadword	IO_RESTART_RSI		
FEB8h	Quadword	IO_RESTART_RDI		
FEC0h	Doubleword	SMMFEC0 [SMM IO Trap Offset]		Read-only
FEC4	Doubleword	SMMFEC4 [Local SMI Status]		Read-only
FEC8h	Byte	SMMFEC8 [SMM IO Restart Byte]		Read-write
FEC9h	Byte	SMMFEC9 [Auto Halt Restart Offset]		Read-write
FECAh	Byte	SMMFECA [NMI Mask]		Read-write
FECBh	5 Bytes	Reserved		
FED0h	Quadword	EFER		Read-only
FED8h	Quadword	SMMFED8 [SMM SVM State]		Read-only
FEE0h	Quadword	Guest VMCB physical address		Read-only
FEE8h	Quadword	SVM Virtual Interrupt Control		Read-only
FEF0h	16 Bytes	Reserved		

Table 10: SMM Save State

Offset	Size	Contents	Access	
FEFCh	Doubleword	SMMFEFC [SMM-Revision Identifier]	Read-only	
FF00h	Doubleword	SMMFF00 [SMM Base Address (SMM_BASE)]	Read-write	
FF04h	28 Bytes	Reserved		
FF20h	Quadword	Guest PAT	Read-only	
FF28h	Quadword	Host EFER ²		
FF30h	Quadword	Host CR4 ²		
FF38h	Quadword	Nested CR3 ²		
FF40h	Quadword	Host Cr0 ²		
FF48h	Quadword	CR4		
FF50h	Quadword	CR3		
FF58h	Quadword	CR0		
FF60h	Quadword	DR7		Read-only
FF68h	Quadword	DR6		
FF70h	Quadword	RFLAGS		Read-write
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		Read-write
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 MSBs of the DW (sometimes called *sign extended*). The 16 LSBs 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
31:23	Reserved.
22	SmiSrcOnLineSpare: SMI source on-line spare. Read-only. This bit is associated with the SMI sources specified in D18F3xB0 [On-Line Spare Control] .
21	Reserved.
20	SmiSrcThrCntL3: SMI source L3 cache thresholding. Read-only. This bit is associated with the SMI source specified in the L3 cache thresholding register (see D18F3x160 [NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)]).
19	SmiSrcThrCntHt: SMI source link thresholding. Read-only. This bit is associated with the SMI source specified in the link thresholding register (see D18F3x160 [NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)]).
18	SmiSrcThrCntDram: SMI source DRAM thresholding. Read-only. This bit is associated with the SMI source specified in the DRAM thresholding register (see D18F3x160 [NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)]).

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][IntPndMsg] (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. <table border="1" data-bbox="256 1150 1453 1392"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>SMM entered from a non-guest state.</td> </tr> <tr> <td>001b</td> <td>Reserved.</td> </tr> <tr> <td>010b</td> <td>SMM entered from a guest state.</td> </tr> <tr> <td>101b-011b</td> <td>Reserved.</td> </tr> <tr> <td>110b</td> <td>SMM entered from a guest state with nested paging enabled.</td> </tr> <tr> <td>111b</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	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.
Bits	Description														
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](#)) are locked from being altered by setting [MSRC001_0015](#)[SmmLock]. SBIOS must 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 [3.10 \[CPUID Instruction Registers\]](#).

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\[LogicalProcessorCount\]](#).

2.4.10.2 L3 Cache Support

The NB includes an L3 cache.

BIOS must take steps to correctly display cache size information on the boot video screen:

- The total cache size is: $(\text{CPUID Fn8000_0005_EDX}[\text{L1IcSize}] * 1 \text{ KB}) + (\text{CPUID Fn8000_0005_ECX}[\text{L1DcSize}] * 1 \text{ KB}) + (\text{CPUID Fn8000_0006_ECX}[\text{L2Size}] * 1 \text{ KB}) + (\text{CPUID Fn8000_0006_EDX}[\text{L3Size}] * 512 \text{ KB})$.
- It is preferred that the BIOS shows the exact breakdown between the L1, L2, and L3 cache sizes and the total. For example, specify L1 (size of IC and DC in KB) + L2 (size of L2 in KB) + L3 (size of L3 in KB) = total cache size in KB.

2.5 Power Management

The processor supports many power management features in a variety of systems. [Table 11](#) provides a summary of ACPI states and power management features and indicates whether they are supported.

Table 11: Power Management Support

ACPI/Power Management State	Supported	Description	
G0/S0/C0: Working	Yes		
G0/S0/C0: Core P-state transitions	Yes	2.5.2 [P-states]	
G0/S0/C0: NB P-state transitions	Yes	2.5.2.2 [NB P-states]	
G0/S0/C0: Hardware thermal control (HTC)	Yes	2.11.2.1 [PROCHOT_L and Hardware Thermal Control (HTC)]	
G0/S0/C0: P-state limit control	Yes	2.11.2.2 [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 [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/C1E: Stop-grant caches not probed	Yes		
G0/S0/Cx: Cache flushing support	Yes		
G1/S1: Stand By (Powered On Suspend)	Yes		
G1/S3: Stand By (Suspend to RAM)	Yes		2.5.5 [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	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 processor electrical data sheet and infrastructure roadmap for power plan definitions. See [1.2 \[Reference Documents\]](#).

The voltage level of VDD and VDDNB may be altered in various states. All the other supplies are fixed.

The processor only supports dual-plane platforms. VDDNB is isolated from VDD on the systemboard and is controlled as a separate voltage plane.

2.5.1.1 Internal VID Registers

The registers within the processor that contain VID fields all use 7-bit VID encodings. See the AMD Voltage Regulator Specification (1.2 [Reference Documents]).

- The VID for VDDNB is controlled by `D18F5x1[6C:60][NbVid]`. See 2.5.2.2.1 [NB P-state Control].
- The VID for VDD is controlled by `MSRC001_00[6B:64][CpuVid]` of the core in the highest-performance P-state.

2.5.1.1.1 VID Encodings

The VID encoding to VDD translations, for both the boot VID and the SVI VID, are defined by the AMD Voltage Regulator Specification (1.2 [Reference Documents]).

The boot VID is 1.0 volts.

2.5.1.1.2 MinVid and MaxVid Check

Hardware limits the minimum and maximum VID code that is sent to the voltage regulator. The allowed limits of MinVid and MaxVid are provided in `MSRC001_0071`. 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.

2.5.1.2 Serial VID Interface

The processor includes an interface, intended to control external voltage regulators, called the serial VID interface (SVI). The SVI encodes voltage regulator control commands, including the VID code, using SMBus protocol over two pins, SVD and SVC, to generate write commands to external voltage regulators. The processor is the master and the voltage regulator(s) are the slave(s). Both pins are outputs of the master; SVD is driven by the slave as well. SVC is a clock that strobes SVD, the data pin, on the rising edge. The frequency of the SVC is controlled by `D18F3xA0[SviHighFreqSel]`. The SVI protocol is specified in the AMD Voltage Regulator Specification (1.2 [Reference Documents]).

2.5.1.3 BIOS Requirements for Power Plane Initialization

- Configure `D18F3xD8[VSRampSlamTime]` based on the platform requirements.
- Configure `D18F3xD4[PowerStepUp, PowerStepDown]`.
- Optionally configure `D18F3xA0[PsiVidEn and PsiVid]`. Refer to 2.5.1.4.1 [PSI_L Bit] for additional details.

2.5.1.4 Low Power Features

2.5.1.4.1 PSI_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. This is supported by the PSI_L bit in the data field of the SVI command. The PSI_L bit is enabled through [D18F3xA0\[PsiVidEn\]](#). The PSI_L bit is asserted if the processor selects a VID code that is higher than or equal to (voltage that is lower than or equal to) the VID code specified in [D18F3xA0\[PsiVid\]](#). PSI_L bit changes only occur on VID changes since it's part of the VID change command.

The voltage regulator must be able to supply the current required for the processor to operate at the VID code specified in [D18F3xA0\[PsiVid\]](#). [MSRC001_00\[6B:64\]\[IddValue, IddDiv\]](#) specify the maximum core current for each P-state. The following algorithm describes how to program PSI_L on the core voltage planes.

```

After cold reset {
  If (PSI is supported on the platform) {
    PSI_vrm_current = current supported by the VDD voltage regulator when PSI is enabled
    PSI_inrush_current = inrush current on the VDD plane during a voltage transition
    Previous_voltage = 7Fh
    For (Pstate_number = 0; Pstate_number <= D18F3xDC\[HwPstateMaxVal\]; Pstate_number++) {
      Pstate_current = ProcIddMax for the Pstate specified by Pstate_number. See
        2.5.2.1.9 \[Processor-Systemboard Power Delivery Compatibility Check\].
      If ((Pstate_number+1) > D18F3xDC\[HwPstateMaxVal\])
        Next_Pstate_current = 0;
      Else Next_Pstate_current = PSI_inrush_current + ProcIddMax for the Pstate specified
        by Pstate_number+1. See 2.5.2.1.9 \[Processor-Systemboard Power Delivery Compatibility Check\].
      Pstate_voltage = MSRC001\_00\[6B:64\]\[CpuVid\] for the Pstate specified by Pstate_number
      If ((Pstate_current <= PSI_vrm_current) && (Next_Pstate_current <= PSI_vrm_current)
        && (Pstate_voltage != Previous_voltage)) {
        Set D18F3xA0\[PsiVid\] = Pstate_voltage
        Set D18F3xA0\[PsiVidEn\] = 1
        Break
      }
      Previous_voltage = Pstate_voltage
    }
  }
}

```

2.5.1.5 Voltage Transitions

The processor supports dynamic voltage transitions on the VDD and VDDNB planes. These transitions are controlled by either hardware or software. VDD and VDDNB voltage levels may be transitioned during state changes involving boot, reset, P-state changes, and C-state changes. In all cases, the voltage is *slammed*; this means that 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 Voltage Regulator Specification, #40182 and [D18F3xD8\[VSRampSlamTime\]](#).

If a voltage increase is requested, the processor waits the amount of time specified by [D18F3xD8\[VSRampSlamTime\]](#) 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 [D18F3xD8\[VSRampSlamTime\]](#) 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.

2.5.1.5.1 Hardware-Initiated Voltage Transitions

When software requests a P-state change using [MSRC001_00\[6B:64\]\[PstateCmd\]](#) hardware coordinates the necessary voltage transitions. See [2.5.2 \[P-states\]](#).

2.5.1.5.2 Software-Initiated Voltage Transitions

The processor supports direct software VID control using [MSRC001_0070 \[COFVID Control\]](#). Hardware P-state transitions using [MSRC001_0062 \[P-state Control\]](#) result in unpredictable behavior if software modifies the NbVid or CpuVid from the appropriate settings for the current P-state reported in [MSRC001_0063 \[P-state Status\]](#).

2.5.1.5.2.1 Software-Initiated NB Voltage Transitions

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

1. Write the destination NbVid to [MSRC001_0070\[NbVid\]](#) on all cores on the node.
2. Wait the specified [D18F3xD8\[VSRampSlamTime\]](#) .

2.5.1.5.2.2 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 P-states

P-states are operational performance states characterized by a unique frequency and voltage. P-states are numbered in ascending order starting with P0. P0 is the highest power, highest performance P-state; each ascending P-state number represents a lower-power, lower-performance P-state than the prior P-state number. As P-state numbers increase, the operating frequency and voltage for a given P-state must be less than or equal to the frequency and voltage of the prior P-state. At least one enabled core (P0) and NB P-state (NB P0) is specified for all processors.

The processor supports independently-controllable frequency planes for each compute unit and the NB; and independently-controllable voltage planes VDD and VDDNB. See [2.5.1 \[Processor Power Planes And Voltage Control\]](#) for voltage plane definitions.

The following terms may be applied to each of these planes:

- FID: frequency ID. Specifies the PLL frequency multiplier, relative to the reference clock, for a given domain.
- DID: divisor ID. Specifies the post-PLL power-of-two divisor that can be used to reduce the operating frequency.
- COF: current operating frequency. Specifies the operating frequency as a function of the FID and DID. Refer to [CoreCOF](#) for the CPU COF formula and [NBCOF](#) for the NB COF formula.
- VID: voltage ID. Specifies the voltage level for a given domain. Refer to [2.5.1.1.1 \[VID Encodings\]](#) for encodings.

All FID and DID parameters for software P-states must be programmed to equivalent values for all cores and NBs in the coherent fabric. See [2.5.2.1.2.1 \[Software P-state Numbering\]](#). Refer to [MSRC001_00\[6B:64\]](#) and

D18F5x1[6C:60] for further details on programming requirements. Processors with different default P-state definitions can be mixed in a multi-socket system and still satisfy the FID and DID programming requirements. See 2.5.2.3 [Mixed Frequency and Power P-state Configuration] for details on multi-socket, mixed-frequency and/or power initialization requirements.

2.5.2.1 Core P-states

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 cores 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]. The FID, DID, and VID for each core P-state is specified in MSRC001_00[6B:64]. The COF for core P-states is a function of half the CLKIN frequency (nominally 100 MHz) and the DID. See MSRC001_00[6B:64][CpuDid] for more details on the DID. Software requests core P-state changes independently using the hardware P-state control mechanism (a.k.a. fire and forget). See 2.5.4 [Frequency and Voltage Domain Dependencies]. Support for hardware P-state control is indicated by CPUID Fn8000_0007_EDX[HwPstate]=1b. P-state transitions using the hardware P-state control mechanism are not allowed until the P-state initialization requirements defined in 2.5.2.1.8 [BIOS Requirements for Core P-state Initialization and Transitions] are complete.

2.5.2.1.1 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 core 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.

- APM consists of CPB and TDP limiting. See 2.5.2.1.1.1 [TDP Limiting].
- APM is enabled if D18F4x15C[ApmMasterEn]=1.
- Support for CPB is specified by CPUID Fn8000_0007_EDX[CPB].
- CPB 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.
- Boosted P-states can be dynamically enabled and disabled through MSRC001_0015[CpbDis]. If core performance boost is disabled, a P-state limit is applied. The P-state limit restricts cores to the highest performance non-boosted P-state.
- APM monitors core activity at the ApmSampleTimer rate, which is specified by D18F4x110[CSampleTimer]. See D18F4x110[CSampleTimer].
- 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.
- Two levels of boosted P-states are supported. Compute units can be placed in the first level of boosted P-states if software requests the highest performance P-state available and processor power consumption remains within the TDP limit. The second level of boosted P-states can only be achieved if a subset of compute units are in CC6, software requests the highest performance P-state available, and the processor power consumption remains within the TDP limit. Compute units that are downcored do not count towards meeting

the [D18F4x16C\[CstateCnt\]](#) requirements. See [D18F4x16C\[CstateCnt, CstateBoost\]](#) and [2.5.2.1.1.1 \[TDP Limiting\]](#).

- 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 [2.5.2.1.11.2 \[_PSS \(Performance Supported States\)\]](#).

2.5.2.1.1.1 TDP Limiting

TDP limiting is a mechanism for capping the power consumption of the processor through a configurable TDP limit. [APM](#) varies the P-state limit of the processor to keep processor power consumption close to the TDP limit within a thermally significant period. See [2.5.2.1.1 \[Application Power Management \(APM\)\]](#). TDP limits can be set by BIOS or APM.

- TDP limiting is enabled if [D18F4x15C\[ApmMasterEn\]=1](#) .
- TDP limits are specified in units of power credits. One TDP unit equals one power credit. TDP units can be converted to watts by multiplying by the scaling factor [MSRC001_0077\[Tdp2Watt\]](#).
- [APM](#) regulates the total processor core TDP. The maximum non-core TDP is specified by [D18F4x1B8\[BaseTdp\]](#). The maximum processor TDP is specified by [D18F4x1B8\[ProcessorTdp\]](#). The maximum total processor core TDP limit can be computed as [D18F4x10C\[NodeTdpLimit\] * \(D18F3xE8\[MultiNodeCpu\] + 1\)](#).
- The minimum total processor core TDP limit is the power consumed when all cores are in [D18F3xDC\[PstateMaxVal\]](#). See [2.5.3.1.1.8 \[Notification of TDP Limit Changes\]](#).
- There are two sources of TDP limits in the processor. [APM](#) uses the lower of the two limits to determine the effective total core TDP limit for the processor.
 - The maximum sum of TDP for all cores on the processor is [D18F4x10C\[NodeTdpLimit\] * \(D18F3xE8\[MultiNodeCpu\] + 1\)](#).
 - An APM defined percentage of the maximum processor TDP limit, [MSRC001_0075\[ApmTdpLimitPercent\]](#).
- The highest performance P-state available to software varies with the TDP limit. See [2.5.2.1.1.2 \[Notification of TDP Limit Changes\]](#) and [2.5.2.1.1 \[Application Power Management \(APM\)\]](#).
- If software requests the highest performance P-state available and processor power consumption remains within the TDP limit, [APM](#) enables transitions to higher performance P-states.
- If processor power consumption exceeds the TDP limit, [APM](#) may restrict transitions to P-states that are lower performance than the software requested P-state.
- The effective total core TDP limit is equally divided between the two nodes on a multi-node processor.

2.5.2.1.1.2 Notification of TDP Limit Changes

APML or BIOS may set a TDP limit that results in limiting the number of P-states available to the OS. When this occurs, BIOS is notified of the new TDP limit via an SMI interrupt. See [D18F4x16C\[ApmTdpLimitIntEn\]](#). In the interrupt routine, BIOS determines the highest performance P-state that satisfies the [D18F5xE8\[ApmTdpLimit\]](#) by performing the following algorithm on all cores of each node. The SMM BIOS should then notify the OS of the new P-state limit via an SCI interrupt.

1. All cores on a node must do the following before any core on the node can perform step 2:
SwPstateRequest = [MSRC001_0062\[PstateCmd\]](#).
2. Only core 0 does the following before any core on the node can perform step 3:
BaseWatt (real number in W) = $((\text{D18F4x1B8[BaseTdp]} * \text{D18F5xE8[Tdp2Watt[15:0]]}) / (2^{16}))$.
ApmWattLimit (real number in W) = $((\text{D18F5xE8[ApmTdpLimit]} * \text{D18F5xE8[Tdp2Watt[15:0]]}) / (2^{16}))$.

```

IF (0 == MSRC001_00[6B:64][PstateEn] indexed by D18F3xDC[HwPstateMaxVal]) THEN
    D18F4x15C[TdpLimitPstate] = MSRC001_0061[PstateMaxVal] - 1.
ELSE
    D18F4x15C[TdpLimitPstate] = MSRC001_0061[PstateMaxVal].
ENDIF.

```

```

For (i=D18F4x15C[NumBoostStates]; i < D18F3xDC[HwPstateMaxVal]; i++) {
PstateWatt[i] (real number in W) = (D18F2xF8[PwrValue[i]] * (D18F5x84[CmpCap] + 1)
    * (D18F3xE8[MultiNodeCpu] + 1)) / 10^D18F2xF8[PwrDiv[i]].
IF ((PstateWatt[i] - BaseWatt) <= ApmWattLimit) THEN
    D18F4x15C[TdpLimitPstate] = i - D18F4x15C[NumBoostStates]
    Break.
ENDIF.
}

```

3. All cores on a node must do the following:
MSRC001_0062[PstateCmd] = SwPstateRequest.

On multi-node processors:

- Both nodes must have the same number of enabled compute units and the number of enabled cores.
- **D18F4x15C**[TdpLimitPstate] must be the same on both nodes.

2.5.2.1.2 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.

2.5.2.1.2.1 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 12: Software P-state Naming

D18F4x15C [NumBoost-States]=1		D18F4x15C [NumBoost-States]=3	
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

Table 12: Software P-state Naming

D18F4x15C[NumBoost-States]=1		D18F4x15C[NumBoost-States]=3	
P-state Name	Corresponding MSR Address	P-state Name	Corresponding MSR Address
P2	MSRC001_0067	P0	MSRC001_0067
P3	MSRC001_0068	P1	MSRC001_0068
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.

2.5.2.1.2.2 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.

2.5.2.1.3 Core P-state Control

Core P-states are dynamically controlled by software and are exposed through ACPI objects (refer to [2.5.2.1.11 \[ACPI Processor P-state Objects\]](#)). 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 [APM](#). See [2.5.2.1.1 \[Application Power Management \(APM\)\]](#).

Hardware sequences the frequency and voltage changes necessary to complete a P-state transition as specified by [2.5.2.1.7 \[Core P-state Transition Behavior\]](#) 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 [2.5.4 \[Frequency and Voltage Domain Dependencies\]](#) for details about hardware coordination.

Table 13: 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

Table 13: 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
P3	3	MSRC001_0068	P1	1	MSRC001_0068
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 2.5.2.1.7 [Core P-state Transition Behavior] for details on hardware P-state voltage control. 2.5.1.3 [BIOS Requirements for Power Plane Initialization] specifies the processor initialization requirements for voltage plane control.

2.5.2.1.4 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 2.5.4 [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 2.5.6 [Effective Frequency Interface].

2.5.2.1.5 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].
- APM. See D18F3xC4 [PstateLimit].
- Core Performance Boost. See 2.5.2.1.1 [Application Power Management (APM)].
- PROCHOT_L assertion. See 2.11.2.1 [PROCHOT_L and Hardware Thermal Control (HTC)].

P-state limits are applied to all cores on the processor. The current P-state limit is provided in MSRC001_0061 [CurPstateLimit]. Changes to the P-state limit 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].

2.5.2.1.6 Core P-state Bandwidth Requirements

- The frequency relationship of (core COF / NB COF) \leq 4 must be maintained for all supported P-state combinations. E.g., a core P0 COF of 4.0 GHz could not be combined with a NB P0 COF of 0.6 GHz; the NB P0 COF would have to be 1.0 GHz or greater; if the NB P0 COF is 1.0 GHz, then the NB P1 COF of 0.5 GHz may only be supported if the corresponding core P-state specify a COF of 2.0 GHz or less.
- All core P-states are required to be defined such that (NB COF/core COF) \leq 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 $\text{CoreCOF} \geq 500 \text{ MHz}$.
- If $(\text{PROC} < \text{OR_C0})$ then all core P-states must be defined such that $\text{CoreCOF} \neq 900 \text{ MHz}$.
- NB COF $\geq 2 * \text{MEMCLK}$ frequency. E.g., for DDR1333 support, NB COF must be 1.4 GHz or higher. (MEMCLK would be 667 MHz).
- NB COF $\geq 800 \text{ MHz}$.
- See 2.12.4 [Link Bandwidth Requirements] for NB COF and link bandwidth requirements.

2.5.2.1.7 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 $\text{D18F3xD8}[\text{VSRampSlamTime}]$ 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 2.5.2.1.10 [BIOS COF and VID Requirements After Warm Reset].
 - After a warm reset $\text{MSRC001_0063}[\text{P-state Status}]$ is consistent with $\text{MSRC001_0071}[\text{CurPstate}]$. $\text{MSRC001_0062}[\text{P-state Control}]$ may not be consistent with $\text{MSRC001_0071}[\text{CurPstate}]$.
 - If $\text{D18F5x1}[\text{6C:60}][\text{NbFid}]$ 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 2.5.1.5.2 [Software-Initiated Voltage Transitions].
- The OS controls the P-state through $\text{MSRC001_0062}[\text{P-state Control}]$, independent of P-state limits described in $\text{D18F3x64}[\text{Hardware Thermal Control (HTC)}][\text{HtcPstateLimit}]$, $\text{D18F3x68}[\text{Software P-state Limit}][\text{StcPstateLimit}]$, and $\text{D18F3xC4}[\text{SBI P-state Limit}]$. 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 $\text{MSRC001_0061}[\text{P-state Current Limit}][\text{PstateMaxVal}]$, regardless of whether it is a limit or OS-requested, then the PstateMaxVal is used instead.

2.5.2.1.8 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, no P-state related ACPI objects should be generated, and BIOS must skip the rest of these steps.
2. Program `D18F3x[84:80] [ACPI Power State Control]` based on BIOS recommendations.
3. Complete the 2.5.1.3 [BIOS Requirements for Power Plane Initialization].
4. Complete the 2.5.2.3.1 [Mixed Northbridge Frequency Configuration Sequence].
5. Transition all parts to the minimum performance P-state using the algorithm detailed in 2.5.2.1.10.2 [Core Minimum P-state Transition Sequence After Warm Reset].
6. Complete the 2.5.2.1.10.3 [NB COF and VID Transition Sequence After Warm Reset]. All cores on a processor must be in the minimum performance P-state prior to executing this sequence.
7. Complete the 2.5.2.1.9 [Processor-Systemboard Power Delivery Compatibility Check].
8. Determine the valid set of P-states:
 - Based on the sequence described in 2.5.2.3 [Mixed Frequency and Power P-state Configuration] for multi-processor systems.
 - Based on the enabled P-states indicated in `MSRC001_00[6B:64] [P-state [7:0]][PstateEn]` for single-processor systems.
9. Transition all cores to the maximum performance P-state by writing 0 to `MSRC001_0062[PstateCmd]`.
10. If P-states are not supported, as indicated by only one enabled selection in `MSRC001_00[6B:64] [P-state [7:0]][PstateEn]`, then BIOS must not generate ACPI-defined P-state objects described in 2.5.2.1.11 [ACPI Processor P-state Objects]. Otherwise, the ACPI objects should be generated to enable P-state support.

2.5.2.1.9 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 should perform this check independently for each node in the coherent fabric. 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. $\text{ProcIddMax} = \text{MSRC001_00[6B:64][IddValue]} * 1/10^{\text{MSRC001_00[6B:64][IddDiv]}} * (\text{CPUID Fn8000_0008_ECX[NC]}+1)$;

The power delivery check should be applied starting with software P0 and continue with increasing P-state indexes (1, 2, 3, and 4) for all enabled P-states. The power delivery check does not need to be applied to boost P-states since APM ensures that the average current consumption over a thermally significant time period remains at or below TDC. Once a compatible software 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`
- `CPUID Fn8000_0008_ECX[NC] = 1d`
- $\text{ProcIddMax} = 32 * 1 * 2 = 64\text{A}$ per plane

The systemboard must be able to supply $\geq 64A$ per plane 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 on a node while performing the power delivery compatibility check then BIOS does not need to take any action.

If at least one P-state is disabled on a node by performing the power delivery compatibility check and at least one P-state remains enabled for that node, 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.
2. If software P0 is disabled, then disable all boosted P-states and program `D18F4x15C[BoostSrc]=0`.
3. Copy the contents of the enabled P-state MSRs (`MSRC001_00[6B: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 2.5.2.1.2.1 [Software P-state Numbering].
4. 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.
5. 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.

1. 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.
2. If `MSRC001_0061[PstateMaxVal] != 000b` copy the contents of the P-state MSR pointed to by `MSRC001_0061[PstateMaxVal]` to the P0 MSR and set `PstateEn=1` for the P0 MSR; 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 2.5.2.1.2.1 [Software P-state Numbering].
3. Adjust the following fields to `D18F4x15C[NumBoostStates]`.
 - `D18F3x64[HtcPstateLimit]`
 - `D18F3x68[SwPstateLimit]`
 - `D18F3xDC[HwPstateMaxVal]`
4. Program `D18F4x15C[BoostSrc]=0`.

2.5.2.1.10 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 2.5.2.1.7 [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 set-

tings 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 2.5.2.1.10.1 [Core Maximum P-state Transition Sequence After Warm Reset] and 2.5.2.1.10.2 [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 2.5.2.1.9 [Processor-Systemboard Power Delivery Compatibility Check] is complete. For processors that support NB P-states, BIOS must also transition NB COF and VID settings after warm reset using the sequence defined in 2.5.2.1.10.3 [NB COF and VID Transition Sequence After Warm Reset].

2.5.2.1.10.1 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]`.

2.5.2.1.10.2 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]` on all cores in the processor. All cores must complete this step before any core can proceed to the next step.
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]`.

2.5.2.1.10.3 NB COF and VID Transition Sequence After Warm Reset

The following sequence must be performed on only one core per node.

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

10. Wait for `D18F5x174[CurNbPstate] = D18F5x170[NbPstateHi]` and `D18F5x174[CurNbFid, CurNbDid]=[NbFid, NbDid]` from `D18F5x1[6C:60]` indexed by `D18F5x170[NbPstateHi]`.
11. Write 0 to `D18F5x170[SwNbPstateLoDis, NbPstateDisOnP0, NbPstateThreshold]`.
12. Wait for `D18F5x174[CurNbPstate] = D18F5x170[NbPstateLo]` and `D18F5x174[CurNbFid, CurNbDid]=[NbFid, NbDid]` from `D18F5x1[6C:60]` indexed by `D18F5x170[NbPstateLo]`.
13. Set `D18F5x170[SwNbPstateLoDis]=Temp1`, `D18F5x170[NbPstateDisOnP0]=Temp2`, and `D18F5x170[NbPstateThreshold]=Temp3`.

2.5.2.1.11 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.

2.5.2.1.11.1 `_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

2.5.2.1.11.2 `_PSS` (Performance Supported States)

A unique `_PSS` entry is created for each non-boosted P-state. See 2.5.2.1.2.1 [Software P-state Numbering]. BIOS must loop through each of `MSRC001_00[6B:64] [P-state [7:0]]` applying the formulas for CoreFreq and Power, and assigning Control and Status appropriately for valid P-states (`PstateEn=1`). BIOS should ignore `MSRC001_00[6B:64] [P-state [7:0]]` for boosted P-states. The `TransitionLatency` and `BusMasterLatency` values can be calculated once for each processor and applied to all `_PSS` entries for cores on that processor.

The Control and Status fields for the highest performance non-boosted P-state must be programmed to 0. Each lower performance non-boosted P-state must have the Control and Status fields programmed in ascending order. The value contained in the 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 in the Control field is a direct indication of the P-state register (`MSRC001_00[6B:64]`) that contains the COF and VID settings for the associated P-state. The value contained in `MSRC001_0063 [P-state Status]` can be used to identify the `_PSS` object of the current P-state by equating `MSRC001_0063[CurPstate]` to the value of the Status field. Refer to 2.5.2 [P-states] for further details on P-state definition and behavior.

- CoreFreq (MHz) = Calculated using the formula for CoreCOF. All CoreFreq values must be rounded to the nearest 100 MHz frequency resulting in a maximum of 50 MHz frequency difference between the reported CoreFreq and calculated CPU COF.
- Power (mW)
 - Convert MSRC001_00[6B:64][CpuVid] to a voltage by referring to the AMD Voltage Regulator Specification (1.2 [Reference Documents]).
 - Power (mW) = voltage * MSRC001_00[6B:64][IddValue] * 1/10^MSRC001_00[6B:64][IddDiv] * 1000
- TransitionLatency (us) and BusMasterLatency (us)
 - If MSRC001_00[6B:64][CpuFid] is the same value for all P-states where MSRC001_00[6B:64][PstateEn]=1: TransitionLatency = BusMasterLatency = (15 steps * D18F3xD4[PowerStepDown] ns/step / 1000 us/ns) + (15 steps * D18F3xD4[PowerStepUp] ns/step / 1000 us/ns)
 - If MSRC001_00[6B:64][CpuFid] is different for any P-states where MSRC001_00[6B:64][PstateEn]=1: TransitionLatency = BusMasterLatency = (15 steps * D18F3xD4[PowerStepDown] ns/step / 1000 us/ns) + D18F3xA0[PIILockTime] us + (15 steps * D18F3xD4[PowerStepUp] ns/step / 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)

2.5.2.1.11.3 _PPC (Performance Present Capabilities)

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

2.5.2.1.11.4 _PSD (P-state Dependency)

For SingleLink processors, the _PSD object is required to be generated for each core as follows:

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

For MultiLink processors, if (DualCoreEnabled==0), then the _PSD object does not need to be generated. If (DualCoreEnabled==1) then the _PSD object must be generated for each core as follows:

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

For SingleLink processors, BIOS should provide an option to choose between either _PSD definition.

2.5.2.1.12 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.

2.5.2.2 NB P-states

The processor supports up to 2 NB P-states (NB P0 through NB P1), specified in [D18F5x1\[6C:60\]](#). Out of cold reset, the voltage and frequency of the NB is specified by [D18F5x174\[StartupNbPstate\]](#).

Support for dynamic NB P-state changes is indicated by [MSRC001_0071\[NbPstateDis\]](#) and more than one enabled selection in [D18F5x1\[6C:60\]\[NbPstateEn\]](#). NB P-state transitions are only supported as part of [2.5.2.3.1 \[Mixed Northbridge Frequency Configuration Sequence\]](#). The FID, DID, and VID for each NB P-state is specified in [D18F5x1\[6C:60\]\[NbFid, NbDid, NbVid\]](#). The COF for NB P-states is a function of the CLKIN frequency (nominally 200MHZ) and [D18F5x1\[6C:60\]\[NbFid, NbDid\]](#).

The two NB P-states are indexed by [D18F5x170\[NbPstateLo, NbPstateHi\]](#) and represent a low and high performance NB P-state. Transitions between [D18F5x170\[NbPstateLo, NbPstateHi\]](#) are dependent on [MSRC001_00\[6B:64\]\[NbPstate\]](#) and [D18F5x170\[NbPstateThreshold\]](#). NB P-states can be disabled by [D18F5x170\[SwNbPstateLoDis, NbPstateDisOnP0\]](#).

2.5.2.2.1 NB P-state Control

The NB P-state is controlled by hardware and is not exposed to dynamic software control through ACPI objects. NB P-state control is specified by [MSRC001_00\[6B:64\]\[NbPstate\]](#).

The Northbridge is placed in [D18F5x170\[NbPstateHi\]](#) if any of the following are true:

- [MSRC001_0071\[NbPstateDis\]](#)=1
- [D18F5x170\[SwNbPstateLoDis\]](#)=1
- [D18F5x170\[NbPstateDisOnP0\]](#)=1 and at least one core on the node is in P0
- The number of compute units in a P-state which specifies [MSRC001_00\[6B:64\]\[NbPstate\]](#)=1 < [D18F5x170\[NbPstateThreshold\]](#)

The Northbridge is placed in [D18F5x170\[NbPstateLo\]](#) if all of the following are true:

- [MSRC001_0071\[NbPstateDis\]](#)=0
- [D18F5x170\[SwNbPstateLoDis\]](#)=0
- [D18F5x170\[NbPstateDisOnP0\]](#)=1 and there are no cores on the node in P0, or [D18F5x170\[NbPstateDisOnP0\]](#)=0
- The number of compute units in a P-state which specifies [MSRC001_00\[6B:64\]\[NbPstate\]](#)=1 >= [D18F5x170\[NbPstateThreshold\]](#)

Changes in the core P-state or [D18F5x170\[SwNbPstateLoDis, NbPstateDisOnP0, NbPstateThreshold\]](#) will cause the NB P-state to be re-evaluated. The current NB P-state is specified by [D18F5x174\[CurNbFid, CurNbDid, CurNbVid\]](#).

2.5.2.3 Mixed Frequency and Power P-state Configuration

Processors with different [CoreCOF](#), [NBCOF](#), and power can be mixed in a system. All NBs must have the same number of P-states. All cores must have the same number of software P-states and all equivalent software P-states must have identical ACPI CoreFreq and Power settings. See [2.5.2.1.2.1 \[Software P-state Numbering\]](#) and [2.5.2.1.11.2 \[_PSS \(Performance Supported States\)\]](#).

- If [MSRC001_00\[6B:64\]\[PstateEn, CpuFid, CpuDid, IddDiv, and IddValue\]](#) are identical for all software P-states on all processors, then no BIOS modifications to [MSRC001_00\[6B:64\]](#) are necessary.

- If `MSRC001_00[6B:64][PstateEn, CpuFid, or CpuDid]` for any software P-state differs between processors, sections 2.5.2.3.3 [Mixed Frequency and Power P-State Configuration Rules] and 2.5.2.3.4 [Mixed Frequency and Power P-State Configuration Sequence] are used to determine the common set of P-states and define the required BIOS modifications to `MSRC001_00[6B:64]`.
- If `MSRC001_00[6B:64][IddDiv or IddValue]` for any software P-state differs between processors and `MSRC001_00[6B:64][PstateEn, CpuFid, and CpuDid]` for all software P-states do not differ between processors, then 2.5.2.3.2 [Mixed Power P-State Configuration Sequence] defines the required BIOS modifications to `MSRC001_00[6B:64]`.
- If `D18F5x1[6C:60][NbFid, NbDid, NbPstateEn]` are identical for all NB P-states on all processors, then no BIOS modifications to `D18F5x1[6C:60]` are necessary. Otherwise, 2.5.2.3.1 [Mixed Northbridge Frequency Configuration Sequence] specifies the required BIOS modifications to `D18F5x1[6C:60]`.
- For all 16-bit coherent links that use 4-bit time operations, if the Northbridge and link frequencies do not meet the requirements in 2.12.4 [Link Bandwidth Requirements], BIOS must modify the link frequencies such that the requirements are met. This should be performed after 2.5.2.3.1 [Mixed Northbridge Frequency Configuration Sequence].

2.5.2.3.1 Mixed Northbridge Frequency Configuration Sequence

BIOS must match `D18F5x1[6C:60][NbFid, NbDid, NbPstateEn]` between all processors in the coherent fabric of a multi-socket system. The lowest setting from all processors in a multi-socket system (determined by using the following equations on each processor and selecting the lowest value) is used as the common `D18F5x1[6C:60][NbFid, NbDid]`. Only processors with the same number of enabled NB P-states (same `D18F5x170[NbPstateMaxVal]`) can be mixed in a system.

```

For each node in the system {
For (i=0; i <= D18F5x170[NbPstateMaxVal]; i++) {
    NewNbFreq = the lowest NBCOF from all processors for NB P-state i
    NewNbFid = D18F5x1[6C:60][NbFid] that corresponds to NewNbFreq
    NewNbDid = D18F5x1[6C:60][NbDid] that corresponds to NewNbFreq
    Write NewNbFid and NewNbDid to D18F5x1[6C:60][NbFid, NbDid] indexed by NB P-state i
    }
IF (D18F5x170[NbPstateMaxVal] == 0) THEN{
    Save D18F5x170 and D18F5x1[6C:60] indexed by NB P-state 1
    Copy D18F5x1[6C:60] indexed by NB P-state 0 to D18F5x1[6C:60] indexed by NB P-state 1
    Write 1 to D18F5x170[NbPstateMaxVal, NbPstateLo]
    Write 0 to D18F5x170[SwNbPstateLoDis, NbPstateDisOnP0, NbPstateThreshold]
    Wait for D18F5x174[CurNbPstate] = D18F5x170[NbPstateLo] and D18F5x174[CurNbFid, CurNb-
Did]=[NbFid, NbDid] from D18F5x1[6C:60] indexed by D18F5x170[NbPstateLo]
    Restore D18F5x170 and D18F5x1[6C:60] indexed by NB P-state 1
    Wait for D18F5x174[CurNbPstate] = D18F5x170[NbPstateHi]
    }
}

```

The new NB frequencies are used in all subsequent NB P-state transitions. See 2.5.2.1.8 [BIOS Requirements for Core P-state Initialization and Transitions].

2.5.2.3.2 Mixed Power P-State Configuration Sequence

BIOS must match `MSRC001_00[6B:64][IddDiv and IddValue]` for each software P-state across all processors using the following sequence. For each `MSRC001_00[6B:64]` with `PstateEn=1`:

1. Read IddDiv and IddValue for all processors.
2. Calculate the resulting power for each processor using the formula documented in 2.5.2.1.11.2 [_PSS (Performance Supported States)].
3. Identify the highest power for all processors.
4. Program IddDiv and IddValue for all processors to the values from the processor with the highest calculated power.

2.5.2.3.3 Mixed Frequency and Power P-State Configuration Rules

- Processors with only one enabled software P-state (`D18F3xDC[HwPstateMaxVal]=D18F4x15C[NumBoostStates]`) cannot be mixed in a system with processors with more than one enabled software P-state (`D18F3xDC[HwPstateMaxVal]! =D18F4x15C[NumBoostStates]`).
- Processors with `D18F3xE8[HtcCapable]=1` cannot be mixed in a system with processors with `D18F3xE8[HtcCapable]=0`.
- In a system where one or more cores are forced down to one software P-state due to board power limitations (see 2.5.2.1.9 [Processor-Systemboard Power Delivery Compatibility Check]), all other cores in the system must be placed into the P-state specified by `D18F3x64[HtcPstateLimit]`. The transition to the HTC P-state can be done at any time during the BIOS POST routine. HTC P-states must be matched according to the guidelines specified in 2.5.2.3.4 [Mixed Frequency and Power P-State Configuration Sequence]. The HTC P-state limit should be used regardless of the `D18F3xE8[HtcCapable]` value. The remaining requirements in this section can be skipped.
- The maximum performance software P0 CoreCOF for the system is equivalent to the lowest software P0 CoreCOF for any processor in the system.
- The number of software P-states for the system is equivalent to, or lower than, the lowest number of software P-states for any processor in the system.
- All CoreCOF calculations are rounded to the nearest 100 MHz frequency for the purposes of frequency matching.
- The CoreCOF for any enabled software P-state can be lowered by modifying `MSRC001_00[6B:64][CpuFid]` from the cold reset value.
- The power for any enabled P-state can be modified by writing to `MSRC001_00[6B:64][IddDiv, IddValue]`.
- P-states can be invalidated by setting `MSRC001_00[6B:64][PstateEn]=0`.
- P-states that are disabled at cold reset should not be enabled.
- `MSRC001_00[6B:64][CpuVid, NbPstate]` cold reset values are not modified by the mixed frequency P-state configuration sequence for any P-state.
- No P-state changes are allowed until all appropriate steps in the sequence have been completed.

2.5.2.3.4 Mixed Frequency and Power P-State Configuration Sequence

All steps in the following sequence use software P-state numbering. See 2.5.2.1.2.1 [Software P-state Numbering].

1. Verify the rules in section 2.5.2.3.3 [Mixed Frequency and Power P-State Configuration Rules] for all processors.
2. Match P0 CPU COF for all cores to the lowest P0 CPU COF value in the coherent fabric, and match P0 power for all cores to the highest P0 power value in the coherent fabric.
 - If all processors have only 1 enabled P-state, the following sequence should be performed on all cores:
 - Write the appropriate CpuFid and CpuDid values resulting from the matched CPU COF to `MSRC001_00[6B:64][CpuFid, CpuDid]` indexed by `D18F3xDC[HwPstateMaxVal]`.
 - Copy `MSRC001_00[6B:64]` indexed by `D18F3xDC[HwPstateMaxVal]` to `MSRC001_00[6B:64]` indexed by `D18F3xDC[HwPstateMaxVal] + 1`.
 - Increment `D18F3xDC[HwPstateMaxVal]` by 1.

- Write 001b to `MSRC001_0062[PstateCmd]`.
- Wait for `MSRC001_0071[CurCpuFid, CurCpuDid] = MSRC001_00[6B:64][CpuFid, CpuDid]` indexed by `D18F3xDC[HwPstateMaxVal]`.
- Write 000b to `MSRC001_0062[PstateCmd]`.
- Wait for `MSRC001_0071[CurPstate] = MSRC001_0071[CurPstateLimit]`.
- Decrement `D18F3xDC[HwPstateMaxVal]` by 1 and exit the sequence (no further steps are required).

Table 14: Representative mixed frequency P-state table example (step 2)

P-state	Cold Reset				Post-Step 2			
	Processor 0	Processor 1	Processor 2	Processor 3	Processor 0	Processor 1	Processor 2	Processor 3
P0	2.5 GHz 90 W	2.7 GHz 90 W	3.2 GHz 100 W	2.5 GHz 70 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W
P1	2.2 GHz ¹ 80 W	2.4 GHz 80 W	3.0 GHz ¹ 90 W	2.3 GHz ¹ 60 W	2.2 GHz ¹ 80 W	2.4 GHz 80 W	3.0 GHz ¹ 90 W	2.3 GHz ¹ 60 W
P2	1.8 GHz 70 W	2.2 GHz ¹ 70 W	2.4 GHz 80 W	1.8 GHz 50 W	1.8 GHz 70 W	2.2 GHz ¹ 70 W	2.4 GHz 80 W	1.8 GHz 50 W
P3	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W
P4	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³

Refer to [Table 19](#) for notes.

- Match the CPU COF and power for P-states used by HTC:
 - Skip to step 4 if any processor reports `D18F3xE8[HTC Capable]=0`.
 - Identify the lowest CPU COF for all processors in the P-state pointed to by `D18F3x64[HtcPstateLimit]`.
 - Modify the CPU COF pointed to by `D18F3x64[HtcPstateLimit]` to the previously identified lowest CPU COF value.
 - Identify the highest power for all processors in the P-state pointed to by `D18F3x64[HtcPstateLimit]`.
 - Modify the power pointed to by `D18F3x64[HtcPstateLimit]` to the previously identified highest power value.

Table 15: Representative mixed frequency P-state table example (step 3)

P-state	Cold Reset				Post-Step 3			
	Processor 0	Processor 1	Processor 2	Processor 3	Processor 0	Processor 1	Processor 2	Processor 3
P0	2.5 GHz 90 W	2.7 GHz 90 W	3.2 GHz 100 W	2.5 GHz 70 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W
P1	2.2 GHz ¹ 80 W	2.4 GHz 80 W	3.0 GHz ¹ 90 W	2.3 GHz ¹ 60 W	2.2 GHz ¹ 90 W	2.4 GHz 80 W	2.2 GHz ¹ 90 W	2.2 GHz ¹ 90 W
P2	1.8 GHz 70 W	2.2 GHz ¹ 70 W	2.4 GHz 80 W	1.8 GHz 50 W	1.8 GHz 70 W	2.2 GHz ¹ 90 W	2.4 GHz 80 W	1.8 GHz 50 W
P3	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W
P4	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³

Refer to [Table 19](#) for notes.

4. Match the CPU COF and power for the lowest performance P-state:
 - If ($D18F3xDC[HwPstateMaxVal] == D18F3x64[HtcPstateLimit]$) for any processor, then set $PstateEn=0$ for all P-states greater than the P-state pointed to by $D18F3x64[HtcPstateLimit]$ for all processors, skip the remaining actions in this step, and go to step 5.
 - Identify the lowest CPU COF for all processors in the P-state pointed to by $D18F3xDC[HwPstateMaxVal]$.
 - Modify the CPU COF for all processors in the P-state pointed to by $D18F3xDC[HwPstateMaxVal]$ to the previously identified lowest CPU COF value.
 - Identify the highest power for all processors in the P-state pointed to by $D18F3xDC[HwPstateMaxVal]$.
 - Modify the power for all processors in the P-state pointed to by $D18F3xDC[HwPstateMaxVal]$ to the previously identified highest power value.

Table 16: Representative mixed frequency P-state table example (step 4)

P-state	Cold Reset				Post-Step 4			
	Processor 0	Processor 1	Processor 2	Processor 3	Processor 0	Processor 1	Processor 2	Processor 3
P0	2.5 GHz 90 W	2.7 GHz 90 W	3.2 GHz 100 W	2.5 GHz 70 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W
P1	2.2 GHz ¹ 80 W	2.4 GHz 80 W	3.0 GHz ¹ 90 W	2.3 GHz ¹ 60 W	2.2 GHz ¹ 90 W	2.4 GHz 80 W	2.2 GHz ¹ 90 W	2.2 GHz ¹ 90 W
P2	1.8 GHz 70 W	2.2 GHz ¹ 70 W	2.4 GHz 80 W	1.8 GHz 50 W	1.8 GHz 70 W	2.2 GHz ¹ 90 W	2.4 GHz 80 W	1.8 GHz 50 W
P3	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 60 W
P4	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³	500 MHz ² 60 W	500 MHz ² 60 W	500 MHz ² 60 W	N/A ³

Refer to [Table 19](#) for notes.

5. Modify $D18F3xDC[HwPstateMaxVal]$ to indicate the lowest performance P-state with $PstateEn$ set for each processor (step 4 can disable P-states pointed to by $D18F3xDC[HwPstateMaxVal]$).
6. Match the CPU COF and power for upper intermediate P-states:
 - Upper intermediate P-states = P-states between (not including) P0 and $D18F3x64[HtcPstateLimit]$.
 - Define each of the available upper intermediate P-states; for each processor concurrently evaluate the following loop; when any processor falls out of the loop (runs out of available upper intermediate P-states) all other processors have their remaining upper intermediate P-states invalidated ($PstateEn=0$); for ($i = D18F3x64[HtcPstateLimit]-1$; $i > D18F4x15C[NumBoostStates]$; $i--$)
 - Identify the lowest CPU COF for P(i).
 - Identify the highest power for P(i).
 - Modify P(i) CPU COF for all processors to the previously identified lowest CPU COF value.
 - Modify P(i) power for all processors to the previously identified highest power value.

Table 17: Representative mixed frequency P-state table example (step 6)

P-state	Cold Reset				Post-Step 6			
	Processor 0	Processor 1	Processor 2	Processor 3	Processor 0	Processor 1	Processor 2	Processor 3
P0	2.5 GHz 90 W	2.7 GHz 90 W	3.2 GHz 100 W	2.5 GHz 70 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W
P1	2.2 GHz ¹ 80 W	2.4 GHz 80 W	3.0 GHz ¹ 90 W	2.3 GHz ¹ 60 W	2.2 GHz ¹ 90 W	N/A ³	2.2 GHz ¹ 90 W	2.2 GHz ¹ 90 W
P2	1.8 GHz 70 W	2.2 GHz ¹ 70 W	2.4 GHz 80 W	1.8 GHz 50 W	1.8 GHz 70 W	2.2 GHz ¹ 90 W	2.4 GHz 80 W	1.8 GHz 50 W
P3	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 60 W
P4	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³	500 MHz ² 60 W	500 MHz ² 60 W	500 MHz ² 60 W	N/A ³

Refer to [Table 19](#) for notes.

Example description:

$D18F3x64[HtcPstateLimit] = D18F4x15C[NumBoostStates] + 1$ for processors 0, 2, and 3. Therefore, processor 1 must have P1 invalidated (remaining upper intermediate P-state). Execution skips to the next numbered step.

7. Match the CPU COF and power for lower intermediate P-states:

- Lower intermediate P-states = P-states between (not including) $D18F3x64[HtcPstateLimit]$ and $D18F3xDC[HwpStateMaxVal]$
- If $D18F3xDC[HwpStateMaxVal] - D18F3x64[HtcPstateLimit] < 2$ for any processor, set $PstateEn=0$ for enabled lower intermediate P-states for all processors with $D18F3xDC[HwpStateMaxVal] - D18F3x64[HtcPstateLimit] > 1$ and skip the remaining actions for this numbered step.
- Define each of the available lower intermediate P-states; for each processor concurrently evaluate the following loop; when any processor falls out of the loop (runs out of available lower intermediate P-states) all other processors have their remaining lower intermediate P-states invalidated ($PstateEn=0$); for ($i = D18F3xDC[HwpStateMaxVal]-1$; $i > D18F3x64[HtcPstateLimit]$; $i--$)
 - Identify the lowest CPU COF for P(i).
 - Identify the highest power P(i).
 - Modify P(i) CPU COF for all processors to the previously identified lowest CPU COF value.
 - Modify P(i) power for all processors to the previously identified highest power value.

Table 18: Representative mixed frequency P-state table example (step 7)

P-state	Cold Reset				Post-Step 7			
	Processor 0	Processor 1	Processor 2	Processor 3	Processor 0	Processor 1	Processor 2	Processor 3
P0	2.5 GHz 90 W	2.7 GHz 90 W	3.2 GHz 100 W	2.5 GHz 70 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W	2.5 GHz 100 W
P1	2.2 GHz ¹ 80 W	2.4 GHz 80 W	3.0 GHz ¹ 90 W	2.3 GHz ¹ 60 W	2.2 GHz ¹ 90 W	N/A ³	2.2 GHz ¹ 90 W	2.2 GHz ¹ 90 W
P2	1.8 GHz 70 W	2.2 GHz ¹ 70 W	2.4 GHz 80 W	1.8 GHz 50 W	N/A ³	2.2 GHz ¹ 90 W	N/A ³	1.0 GHz 70 W
P3	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W	1.0 GHz 70 W	1.0 GHz 70 W	1.0 GHz 70 W	500 MHz ² 60 W
P4	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³	500 MHz ² 60 W	500 MHz ² 60 W	500 MHz ² 60 W	N/A ³

Refer to [Table 19](#) for notes.

Example description:

$D18F3xDC[HwPstateMaxVal] - D18F3x64[HtcPstateLimit] > 1$ for all processors. Therefore, the conditions of the first bullet are not met and execution continues to bullet two.

Loop index *i* initializes to:

P3 for processors 0, 1, and 2

P2 for processor 3

On the first iteration of the loop processor 1 has the lowest CPU COF of 1.0 GHz for P(*i*), and processor 2 has the highest power of 70 W for P(*i*). The P(*i*) values of each processor are modified to 1.0 GHz and 70 W.

The loop index *i* is decremented for all processors.

Processor 1 fails the loop index test of $i > D18F3x64[HtcPstateLimit]$ with $i = 2$ and $D18F3x64[HtcPstateLimit] = 2$. Processor 3 fails the loop index test of $i > D18F3x64[HtcPstateLimit]$ with $i = 1$ and $D18F3x64[HtcPstateLimit] = 1$.

P2 is invalidated for processors 0 and 2 (remaining lower intermediate P-state).

Execution skips to the next numbered step.

8. Place all cores into a valid COF and VID configuration corresponding to an enabled P-state:
 - Select an enabled P-state not equal to the P-state pointed to by $MSRC001_0063[CurPstate]$ for each core.
 - Transition all cores to the selected P-states by writing the Control value from the $_PSS$ object corresponding to the selected P-state to $MSRC001_0062[PstateCmd]$.
 - Wait for all cores to report the Status value from the $_PSS$ object corresponding to the selected P-state in $MSRC001_0063[CurPstate]$.

Table 19: Representative mixed frequency P-state table example (final)

P-state	Cold Reset				Post-Algorithm			
	Processor 0	Processor 1	Processor 2	Processor 3	Processor 0	Processor 1	Processor 2	Processor 3
P0	2.5 GHz 90 W	2.7 GHz 90 W	3.2 GHz 100 W	2.5 GHz 70 W	2.5 GHz <i>100 W</i>	2.5 GHz <i>100 W</i>	2.5 GHz 100 W	2.5 GHz <i>100 W</i>
P1	2.2 GHz ¹ 80 W	2.4 GHz 80 W	3.0 GHz ¹ 90 W	2.3 GHz ¹ 60 W	2.2 GHz ¹ 90 W	N/A ³	2.2 GHz ¹ 90 W	2.2 GHz ¹ 90 W
P2	1.8 GHz 70 W	2.2 GHz ¹ 70 W	2.4 GHz 80 W	1.8 GHz 50 W	N/A ³	2.2 GHz ¹ 90 W	N/A ³	1.0 GHz 70 W
P3	1.2 GHz 60 W	1.0 GHz 60 W	1.6 GHz 70 W	500 MHz ² 40 W	1.0 GHz 70 W	1.0 GHz 70 W	1.0 GHz 70 W	500 MHz ² 60 W
P4	500 MHz ² 50 W	600 MHz ² 50 W	500 MHz ² 60 W	N/A ³	500 MHz ² 60 W	500 MHz ² 60 W	500 MHz ² 60 W	N/A ³

Notes:

- 1) Indicates the P-state pointed to by `D18F3x64[HtcPstateLimit]`.
 - 2) Indicates the P-state pointed to by `D18F3xDC[HwPstateMaxVal]`.
 - 3) N/A indicates a P-state with `MSRC001_00[6B:64][PstateEn]=0`.
- *Italics* indicates values modified by the mixed frequency P-state algorithm from the cold reset value.
 - ***Bold Italics*** indicates values modified by this step of the mixed frequency P-state algorithm from the cold reset value.
 - `MSRC001_00[6B:64][CpuVid, NbPstate]` are not modified by the mixed frequency P-state algorithm and are not shown.

Table 20: Representative mixed frequency _PSS object example

P-state	Post-Algorithm _PSS			
	Processor 0	Processor 1	Processor 2	Processor 3
0	CoreFreq = 2.5 GHz Power = 100 W Control = Status = 0h	CoreFreq = 2.5 GHz Power = 100 W Control = Status = 0h	CoreFreq = 2.5 GHz Power = 100 W Control = Status = 0h	CoreFreq = 2.5 GHz Power = 100 W Control = Status = 0h
1	CoreFreq = 2.2 GHz* Power = 90 W Control = Status = 1h	CoreFreq = 2.2 GHz* Power = 90 W Control = Status = 2h	CoreFreq = 2.2 GHz* Power = 90 W Control = Status = 1h	CoreFreq = 2.2 GHz* Power = 90 W Control = Status = 1h
2	CoreFreq = 1.0 GHz Power = 70 W Control = Status = 3h	CoreFreq = 1.0 GHz Power = 70 W Control = Status = 3h	CoreFreq = 1.0 GHz Power = 70 W Control = Status = 3h	CoreFreq = 1.0 GHz Power = 70 W Control = Status = 2h
3	CoreFreq = 500 MHz Power = 60 W Control = Status = 4h	CoreFreq = 500 MHz Power = 60 W Control = Status = 4h	CoreFreq = 500 MHz Power = 60 W Control = Status = 4h	CoreFreq = 500 MHz Power = 60 W Control = Status = 3h

Notes:

- * Indicates the P-state pointed to by `D18F3x64[HtcPstateLimit]` at cold reset.
- Refer to [2.5.2.1.11.2 \[_PSS \(Performance Supported States\)\]](#) for details on _PSS object creation and field definitions for CoreFreq, Power, Control, and Status.
- TransitionLatency and BusMasterLatency are not modified by the mixed frequency P-state algorithm and are not shown.
- Units are not indicative of the conventions required by the ACPI _PSS object. Refer to [2.5.2.1.11.2 \[_PSS](#)

([Performance Supported States](#)) for details on `_PSS` object creation.

2.5.3 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.

2.5.3.1 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 per-core IO-based C-states. Up to three per-core IO-based C-states are supported, per-core IO-based C-state 0, 1, and 2. The per-core 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 \[C-state Request Interface\]](#). The actions taken by the processor when entering a low-power C-state are configured by software. See [2.5.3.3 \[C-state Actions\]](#) for information about AMD specific actions.

2.5.3.2 C-state Request Interface

C-states are dynamically requested by software and are exposed through ACPI objects (see [2.5.3.6 \[ACPI Processor C-state Objects\]](#)). C-states can be requested on a per-core basis. Software requests a C-state change in one of two ways, either by executing the HLT instruction or by reading from an IO address specified by `MSRC001_0073[CstateAddr]` plus an offset of 0 through 7 (see `D18F4x11[C:8]`). The processor always returns 0 for this IO read. Offsets 2 through 7 result in an offset of 2. Executing the HLT instruction 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 [2.5.4 \[Frequency and Voltage Domain Dependencies\]](#) and [2.5.3.3 \[C-state Actions\]](#).

2.5.3.3 C-state Actions

A core takes one of several different possible actions based upon a C-state change request from software. Regardless of the method used to enter a C-state (see [2.5.3.2 \[C-state Request Interface\]](#)), `D18F4x128[CoreCstateMode]` indicates whether the C-state actions are specified by the C1 SMAF code or the C-state action field that corresponds to the IO address read by software. The C1 SMAF code is defined in `F3x84[31:24]` and the C-state action fields are defined in `D18F4x11[C:8]`.

2.5.3.3.1 C-state Divisors and Probes

When a core enters a non-C0 state, its clock is ramped down to a divisor specified by `D18F3x[84:80][ClkDivisor]` or `D18F4x11[C:8][ClkDivisor]`. If probes occur during this time 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][CpuPrbEn]`. See [2.5.3.3.2 \[C-state Cache Flush\]](#).

2.5.3.3.2 C-state Cache Flush

A core can flush its L1 and L2 caches after it enters any non-C0 state. Once a core flushes its caches, probes are no longer sent to that core. This improves probing performance for cores that are in C0.

If a core enters a non-C0 state and cache flush is enabled a timer counts down for a programmable period of time.

- Cache flush is enabled if $((D18F4x128[CoreCstateMode]==1) ? (D18F3xDC[CacheFlushOnHaltCtl]!=0) : (D18F4x11[C:8][CacheFlushEn]==1))$
- The programmable period of time is specified by $D18F3xDC[CacheFlushOnHaltTmr]$ if $(D18F4x128[CoreCstateMode]==1)$ or $D18F4x11[C:8][CacheFlushTmrSel]$ if $(D18F4x128[CoreCstateMode]==0)$.
- Each compute unit has one timer that is shared by both cores.

When the timer expires, the core flushes its L1 and L2 caches to either the L3 (if the processor supports an L3) or to DRAM (if the processor does not support an L3). During the cache flush, clocks are ramped down as specified by $D18F3x[84:80][ClkDivisor]$ if $(D18F4x128[CoreCstateMode]==0)$ or by $D18F4x11[C:8][ClkDivisor]$ if $(D18F4x128[CoreCstateMode]==1)$. After the cache flush is complete, clocks are ramped down as specified by $D18F3xDC[CacheFlushOnHaltCtl]$.

The timer is reset if the core exits the C-state for any reason. 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.

2.5.3.3.3 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. Power gating reduces the amount of dynamic power consumed by the core. VDD is not reduced when a core is in CC6. The following actions are taken by hardware prior to CC6 entry:

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 2.5.3.3.2 [C-state Cache Flush].
3. Internal core state is saved to DRAM.
4. Power is removed from the core.

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

1. $D18F4x11[C:8][CacheFlushEn]=1$ for the corresponding C-state action field.
2. $D18F4x11[C:8][CacheFlushTmrSel] \neq 11b$ for the corresponding C-state action field.
3. $D18F4x11[C:8][PwrGateEn]=1$ for the corresponding C-state action field.
4. $D18F2x118[CC6SaveEn]=1$.
5. $D18F2x118[LockDramCfg]=1$.
6. $D18F4x128[CoreCstateMode]=0$.
7. The CC6 storage area in DRAM is configured. See 2.10.8 [DRAM CC6 Storage].

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.

2.5.3.4 Exiting C-states

The following events may cause the processor to exit from a HLT initiated C-state:

- INTR
- NMI
- SMI
- INIT
- STARTUP

The following events always cause the processor to exit from a HLT initiated C-state:

- STPCLK assertion messages.
- RESET_L assertion.

The following events always cause the processor to exit from an IO read initiated C-state:

- INTR
- NMI
- SMI
- INIT
- STARTUP
- STPCLK assertion messages
- RESET_L assertion

The following events never cause the processor to exit from a HLT or IO read initiated C-state:

- A20M
- IGNNE
- machine check error

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.

2.5.3.5 C1 Enhanced State (C1E)

The C1 enhanced state (C1E) is a stop-grant state supported by the processor. The C1E state is characterized by the following properties:

- All cores are in a non-C0 state.
- The ACPI-defined P_LVL3 register has been accessed.
- The chipset has issued a STPCLK assertion message with the appropriate SMAF for C1E entry. Note that [D18F3x\[84:80\]](#) specify the processor clocking and voltage behavior in response to the C1E SMAF.
- The processor has issued a STOP_GRANT message to the chipset.

General requirements for C1E:

- The ACPI-defined C2 and C3 states that require IO hub support must not be declared to the operating system.
- C1E should only be enabled when the platform is in ACPI power management mode.

2.5.3.5.1 Message Triggered C1E

Message triggered C1E is built upon dedicated hardware, a set of specific messages, board-level sideband signals, and timers in the IO hub.

The following board-level link-defined sideband signals are used in this protocol:

- ALLOW_LDTSTOP: Board-level wired-OR signal. ALLOW_LDTSTOP is driven by all chipset NBs in the system and is an input to the IO hub. When ALLOW_LDTSTOP is asserted, the IO hub is permitted

to assert LDTSTOP_L. When ALLOW_LDTSTOP is deasserted the IO hub must deassert LDTSTOP_L, and keep it deasserted until ALLOW_LDTSTOP asserts, after which a programmable period of time must pass before LDTSTOP_L can be asserted. The IO hub must be configured to treat a deassertion of ALLOW_LDTSTOP as also being an assertion of BMREQ#. BMREQ# is the bus-mastering request input of the IO hub, and is a distinct input from ALLOW_LDTSTOP.

- IDLE_EXIT_L: Board-level wired-OR signal. IDLE_EXIT_L is driven by all processors in the system and is an input to the IO hub. IDLE_EXIT_L is asserted by the processor when it receives a Fixed or ExtInt interrupt while the target core is in STOP_GRANT. When IDLE_EXIT_L is asserted, the IO hub initiates the C1E exit sequence. IDLE_EXIT_L is deasserted by the processor when the target core exits STOP_GRANT. See also [D18F3xA0\[IdleExitEn\]](#).

2.5.3.5.2 BIOS Requirements to Initialize Message Triggered C1E

To enable and configure message triggered C1E, software must program the following registers on all nodes in the system:

- [D18F3xD4\[MTC1eEn\]](#)=1.
- [MSRC001_0055\[BmStsClrOnHaltEn\]](#)=1; [MSRC001_0055\[IOMsgAddr\]](#)="the base address of the BM_STS register in the IO hub".
- See [D18F3xD4\[CacheFlushImmOnAllHalt\]](#).
- See [D18F3xDC\[CacheFlushOnHaltCtl\]](#).
- See [D18F4x11\[C:8\]\[CacheFlushEnCstAct0, CacheFlushEnCstAct1\]](#).
- See [D18F3xD4\[StutterScrubEn\]](#).
- [MSRC001_0015\[HltXSpCycEn\]](#)= 1.
- [D18F3xA0\[IdleExitEn\]](#) = 1.
- [D18F4x128\[CstateMsgDis\]](#) = 0.

2.5.3.6 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 [2.5.3.6.1 \[_CST\]](#) and [2.5.3.6.2 \[_CSD\]](#).

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 [2.5.3.6.4 \[Fixed ACPI Description Table \(FADT\) Entries\]](#).

2.5.3.6.1 _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.

2.5.3.6.2 _CSD

If ([DualCoreEnabled](#)==1) then 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.

If (`DualCoreEnabled==0`) then the `_CSD` object should not be generated.

2.5.3.6.3 `_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.

2.5.3.6.4 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.

BIOS must declare the `PSTATE_CNT` entry as 00h.

2.5.3.7 BIOS Requirements for Initialization

1. Initialize `MSRC001_0073[CstateAddr]` with an available IO address. See 2.5.3.6.4 [Fixed ACPI Description Table (FADT) Entries].
2. Initialize `D18F4x11[C:8]`.
3. Generate ACPI objects as described in 2.5.3.6 [ACPI Processor C-state Objects].

2.5.4 Frequency and Voltage Domain Dependencies

Whenever a P-state or C-state is requested on a given core, hardware must take frequency and voltage domain dependencies into account when deciding whether to make the requested state change. Cores within a compute unit share a common frequency and voltage domain, so they must always be at the same P-state and C-state. Compute units within a processor have independent frequency domains, but share a common voltage domain. The voltage is determined by the highest-performance P-state requested on any core.

2.5.4.1 Dependencies Between Cores in a Compute Unit

When software requests different power management states for the cores in a compute unit, hardware determines which state to target as follows:

Table 21: Dependencies Between Cores in a Compute Unit

Core 0 Request	Core 1 Request	Policy	Compute Unit State
C-state	C-state	Frequency is determined by D18F4x128[CoreCstatePolicy] . Voltage is determined by D18F4x128[CoreCstatePolicy] if there is a voltage change associated with the C-state, otherwise voltage is determined by the highest performance P-state requested on any core.	Highest performance P-state, C-state is determined by D18F4x128[CoreC-statePolicy]
P-state	P-state	Frequency and voltage is determined by the highest performance P-state requested on any core	Highest performance P-state, C0
C-state	P-state	Frequency and voltage is determined by the highest performance P-state requested on any core	Highest performance P-state, C0
P-state	C-state	Frequency and voltage is determined by the highest performance P-state requested on any core	Highest performance P-state, C0

2.5.5 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.

2.5.5.1 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, system memory enters self-refresh mode. Software is responsible for bringing memory out of self-refresh mode when resuming from S3. To bring memory out of self-refresh mode, see [2.10.5 \[DCT/DRAM Initialization and Resume\]](#).

Many of the systemboard power planes for the processor are powered down during S3. Refer to the electrical data sheet for the following ([1.2 \[Reference Documents\]](#)):

- 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.

2.5.6 Effective Frequency Interface

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 by periodically reading two counters, APERF ([MSR0000_00E8 \[Actual Performance Frequency Clock Count \(APERF\)\]](#)) and MPERF ([MSR0000_00E7](#)

[Max Performance Frequency Clock Count (MPERF)].

Effective frequency is calculated as follows:

1. Read APERF and MPERF.
2. An interval greater than 1 ms.
3. Read APERF and MPERF.
4. Calculate effective frequency = $((\text{APERF}_{\text{NEW}} - \text{APERF}_{\text{OLD}}) / \{\text{MPERF}_{\text{NEW}} - \text{MPERF}_{\text{OLD}}\}) * \text{P0 frequency}$ using software P-state numbering.

Additional notes:

- When reading or writing APERF and MPERF, software must disable interrupts or any other events that may occur in between writing APERF and writing MPERF or between reading APERF and reading MPERF.
- The behavior of APERF and MPERF may be modified by [MSRC001_0015](#)[EffFreqCntMwait].
- The accuracy of the effective frequency interface is +/- 50MHz if the following constraints are met:
 - Effective frequency is calculated no more frequently than once per 1 ms period.
 - Software executes only MOV instructions, and no more than 3 MOV instructions, between reading or writing APERF and MPERF.
 - The other core of the compute unit is not performing a long latency operation. For example: cache flush operations such as the WBINVD instruction, IO read operations, and switching between host and SVM guest modes.
- APERF and MPERF are invalid if an overflow occurs.

2.6 Processor State Transition Sequences

2.6.1 ACPI Power State Transitions

This section specifies ACPI power state transitions as controlled by [D18F3x\[84:80\]](#) or [D18F4x11\[C:8\]](#). All references to [D18F4x11\[C:8\]](#) in this section also apply to [D18F3x\[84:80\]](#). See [D18F4x128](#)[CoreCstateMode].

The following describes the state transition behavior associated with ACPI power state transitions:

- All C-state controllable parameters take effect after an LDTSTOP_L assertion except ClkDivisorCstAct, which takes effect before the LDTSTOP_L assertion.
- ClkDivisorCstAct:
 - [D18F4x11\[C:8\]](#)[ClkDivisorCstAct] is applied after the processor has transitioned from C0 to a low-power state (halt or stop-grant) and the hysteresis time ([D18F3xD4](#)[ClkRampHystSel]) has elapsed.
 - [D18F4x11\[C:8\]](#)[ClkDivisorCstAct] is removed when the processor transitions from a low-power state to C0 or when a probe occurs and [D18F4x11\[C:8\]](#)[CpuPrbEnCstAct]=0.
- Probes:
 - [D18F4x11\[C:8\]](#)[CpuPrbEnCstAct] specifies how probes are handled while in the low-power state.
- DRAM:
 - DRAM self refresh is enabled and the memory clock is tristated if LDTSTOP_L is asserted.
 - When LDTSTOP_L is deasserted the DRAM memory clock is enabled and self refresh is disabled.
 - Occurs in parallel to re-connecting the link.

Refer to the HyperTransport™ link specification for system management sequencing requirements when performing ACPI state transitions.

2.7 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: [2.7.1](#) [Core Per-

formance Monitor Counters], consisting of one set located in each core of each compute unit and 2.7.2 [NB Performance Monitor Counters].

2.7.1 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 3.15 [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 22; 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 22: 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 22: 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.

2.7.2 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 [3.16 \[NB Performance Counter Events\]](#).

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.

2.7.3 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](#).

2.8 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 [2.8.3 \[Processor Configuration Space\]](#).

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 \[NB 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]}.

The 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.

2.8.1 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 use 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.

2.8.2 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.

2.8.3 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.

2.9 Northbridge (NB)

Each node includes a single northbridge that provides the interface to the local core(s), the interface to system memory, the interface to other processors, 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 of each node is responsible for routing transactions sourced from cores and links to the appropriate core, cache, DRAM, or link. See 2.4.5 [System Address Map].

2.9.1 NB Architecture

Major NB blocks are: System Request Interface (SRI), Memory Controller (MCT), DRAM Controllers (DCTs), L3 cache, 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 links.

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[17C:140,7C:40] [DRAM Base/Limit].

2.9.2 GART

Note: The GART has been deprecated for future use and new product development; it is supported only for compatibility with existing products.

The GART is a device that translates a range of physical address space, called the GART aperture, to a logical address based on page tables in system memory. The GART also includes a 16 entry cache for the page table translations. The registers that specify GART behavior are:

- D18F3x90 [GART Aperture Control].
- D18F3x94 [GART Aperture Base]
- D18F3x98 [GART Table Base]
- D18F3x9C [GART Cache Control]

Programming requirements:

- GART translations to addresses above 1 terabyte are not supported.
- The page table is required to reside within DRAM.
- The page table is required to be mapped to the UC memory type or be updated with strongly ordered UC stores.
- D18F3x90[DisGartTblWlkPrb] must be set to 1.
- The page tables are expected to translate to DRAM address ranges only; translations to MMIO ranges result in undefined behavior.

- The GART registers must be programmed to the same value for all nodes in the system.

2.9.3 NB Routing

There are two factors that affect how a transaction is routed:

1. Address space routing determines the destination node.
2. HyperTransport™ transaction routing determines the path in the coherent fabric that the transaction follows to reach the destination node.

2.9.3.1 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.

2.9.3.1.1 DRAM and MMIO Memory Space

Core 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:

- If the physical address matches against the AGP-aperture (D18F3x90, D18F3x94), then the physical address is translated through the AGP GART. The AGP aperture match ignores the access DRAM/MMIO destination and overrides any match to D18F1x[17C:140,7C:40] [DRAM Base/Limit], and D18F1x[1CC:180,BC:80] [MMIO Base/Limit].
 - For accesses from IO devices, the cacheability attribute from the GART entry's "Coherent" bit, as specified in D18F3x98[GartTblBaseAddr], is applied.
 - Upstream IO accesses to the GART aperture are affected by MSRC001_0015[ForceRdWrSzPrb].
 - For accesses from a CPU, the attribute already applied by the core is used and the GART entry's "Coherent" bit is ignored.
 - System software should ensure that the cacheability attribute assigned to an AGP aperture matches the "Coherent" bit in the matching GART entry.
- IO-device accesses that do not match the AGP aperture and post-GART translated addresses are compared against:
 - If the access matches D18F1x[1CC:180,BC:80] [MMIO Base/Limit], then the transaction is routed to the specified link;
 - Else, if the access matches D18F1x[17C:140,7C:40] [DRAM Base/Limit], then the access is routed to the specified link or DCT;
 - Else, the access is routed to the node or link that contains compatibility (subtractive) address space, specified by D18F0x60[SbNode] and D18F0x64[SbLink].
- For core accesses the routing is determined based on the DRAM/MMIO destination:
 - If the destination is DRAM:
 - If the access matches D18F1x[17C:140,7C:40] [DRAM Base/Limit], 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].
 - If the destination is MMIO:
 - 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, If the access matches [D18F1x\[1CC:180,BC:80\] \[MMIO Base/Limit\]](#), 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\[SbLink\]](#).

2.9.3.1.2 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 specified link;
- 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 node or link that contains compatibility (subtractive) address space, specified by [D18F0x60\[SbNode\]](#) and [D18F0x64\[SbLink\]](#).

2.9.3.1.3 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\[SbLink\]](#).

2.9.3.2 HyperTransport™ Technology Routing

There are three types of HyperTransport™ transactions routed by the NB: (1) broadcast transactions, (2) request transactions, and (3) response transactions. The NB includes routing registers for each node that specify the link to route each transaction type accessed through [D18F0x\[5C:40\] \[Routing Table\]](#).

[Figure 5](#) through [Figure 7](#) show the supported topologies for G34 processors. The links between internal nodes are described in [2.12.1.5 \[Link Mapping Between Package and Node\]](#). The IO hub must be connected to internal node 0.

[Figure 5](#) shows the supported two processor topology for G34r1 processors. For the link connections between the two processors, the thick lines represent 16-bit links and the thin lines represent 8-bit links. Optionally, the 16-bit links may be swapped with the 8 bit links. The links between internal nodes are described in [2.12.1.5 \[Link Mapping Between Package and Node\]](#).




Figure 5: Two Processor G34r1 Topology

Figure 6 shows the supported four processor topology for G34r1 processors with two 16-bit IO links. This topology provides maximum performance. Figure 7 shows the supported four processor topology for G34r1 processors with four 16-bit IO links. This topology provides maximum IO. The links between processors are all 8-bits.




Figure 6: Four Processor G34r1 Maximum Performance Topology




Figure 7: Four Processor G34r1 Maximum IO Topology

2.9.3.2.1 Routing Table Configuration

The routing table registers must be configured correctly in multi-node systems to ensure that probes are only delivered once to each node and to ensure that the routing table is deadlock free.

A routing table is deadlock free if it contains no open-paths and no two-hop cycles.

An open-path is a routing path between nodes that traverse one or more nodes that contains a subpath that is not a routing path in the routing table. For example if the routing path between nodes 0 and 2 in Figure 8 was Node 0->Node 1->Node 3->Node 2 and the routing path between Nodes 3 and 2 was not Node 3->Node 2 then the routing path between Nodes 0 and 2 would be open because the subpath Node 3->Node 2 is not a path in the routing table.




Figure 8: Sample Four-Node Configuration

A two-hop cycle is a group of two hop routing paths (routing paths between two nodes that pass through a third node) such that the first and second nodes in the each two hop routing path are also the second and third nodes in a two hop routing path in the group.

Consider the four node configuration shown in Figure 8. A two-hop cycle would occur in this configuration if the routing table was configured with the following routing paths:

- The routing path from Node 0 to Node 3 is: Node 0->Node 1->Node 3.

- The routing path from Node 1 to Node 2 is: Node 1->Node 3->Node 2.
- The routing path from Node 2 to Node 1 is: Node 2->Node 0->Node 1.
- The routing path from Node 3 to Node 0 is: Node 3->Node 2->Node 0.

To break this cycle at least one but no more than three of these routing paths must be modified to use a different intermediate node. Reconfiguring the routing paths as follows eliminates the 2-hop cycle.

- The routing path from Node 0 to Node 3 is: Node 0->Node 1->Node 3.
- The routing path from Node 1 to Node 2 is: Node 1->Node 3 to Node 2.
- The routing path from Node 2 to Node 1 is: Node 2->Node 0->Node 1.
- The routing path from Node 3 to Node 0 is: Node 3->Node 1->Node 0.

2.9.3.2.2 NodeId Enumeration Requirements for Dual-node Processors

The assignment of [D18F0x60](#)[NodeId] in a dual-node processor system, identified by [D18F3xE8](#)[MultiNodeCpu], is performed starting with the node containing the BSC. The BSC node by definition has NodeId=0. The other internal node in that dual-node processor is assigned NodeId=1. See [D18F0x1A0](#)[IntLnkRoute] for how to identify the internal nodes in a dual-node processor. The remaining NodeId's in the system must be assigned using the following rules:

- Internal nodes must be assigned NodeId's in even/odd pairs. A pair of NodeId's consists of N and N+1 where N is an even integer.
- NodeId's must be assigned contiguously. For example, node ID assignment {0, 1, 2, 3} is allowed but {0, 1, 4, 5} is not.

For example, consider a 3 processor system consisting of dual-node processors A, B, and C where internal node 1 of processor B contains the BSC. A valid NodeId assignment would be:

NodeId	Processor	D18F3xE8 [IntNodeNum]
0	B	1
1	B	0
2	A	0
3	A	1
4	C	1
5	C	0

2.9.3.2.3 BIOS Requirements for Systems with Mixed Processor Families

Processors that are not covered by this document are not supported on coherent links by processors covered by this document. BIOS must ensure that all nodes in the coherent fabric are processors covered by this document by reading [D18F3xFC](#) [CPUID Family/Model/Stepping] before initializing the node. If a node that is not a processor covered by this document is discovered, BIOS must configure the BSP routing tables as a single processor system.

The BIOS may continue the boot process in order to display an error message on the screen if the BSP has DRAM attached and the display adapter is connected to an IO link accessible to the BSP. If these conditions are not met the BIOS may signal an error in a implementation specific manner. The BIOS must not continue the boot process after the error has been reported.

2.9.3.2.4 Link Traffic Distribution

Link traffic distribution is a mechanism to reduce coherent link congestion by distributing the traffic over mul-

multiple links. There are two supported modes for link traffic distribution and BIOS must ensure that only one mode is enabled in a system. See 2.9.3.2.4.1 [Coherent Link Traffic Distribution] and 2.9.3.2.4.2 [Coherent Link Pair Traffic Distribution].

See 2.9.3.2.4.3 [Victim Distribution Mode].

2.9.3.2.4.1 Coherent Link Traffic Distribution

Coherent link traffic distribution supports a 2-node systems in which multiple coherent links (2 or 3) are connected between the nodes in order to increase bandwidth between them. All links between the two nodes should be the same width (either 16-bit ganged links or 8-bit unganged sublinks). This mode is enabled by `D18F0x164` [Coherent Link Traffic Distribution]. The following requirement must be met:

- For all virtual channels that are enabled for distribution, the corresponding routing table entry in `D18F0x[5C:40]` is required to select one of the links specified for distribution in `D18F0x164`[DstLnk].

2.9.3.2.4.2 Coherent Link Pair Traffic Distribution

Coherent link pair traffic distribution supports up to four coherent link groupings among two or more nodes in a system. This mode is enabled by `D18F0x1E0` [Coherent Link Pair Traffic Distribution]. Each link group consists of a pair of coherent links between two nodes. One of the links in the link pair is designated as the master link and the other link is designated as the alternate link. The coherent link pair can be either symmetric (same width) or asymmetric (different width). The following requirements must be met:

- For an asymmetric link pair, the master link must be the larger link.
- The master link is specified and the alternate link is not specified in `D18F0x[5C:40]`.

2.9.3.2.4.3 Victim Distribution Mode

Victim distribution mode is a way to direct victim traffic on to ganged links and away from unganged links as a way to reduce unganged link congestion.

Victim distribution mode is enabled for a system only if all of the following are true:

- 2 processor (4 node) G34 system.

Victim distribution mode is enabled on a node only if all of the following are true:

- Victim distribution mode is enabled for the system.
- The node connects to another node directly with only 1 unganged link hop and indirectly through 2 ganged link hops.

A node that enables victim distribution mode is programmed as follows:

- `D18F0x164`[cHTVicDistMode]=1.
- `D18F0x164`[DstNode] is set to the node that is 2 ganged link hops from this node. Victims not to this node will be directed as normal.
- `D18F0x164`[cHTPrbDistEn]=0; `D18F0x164`[cHTRspDistEn]=1; `D18F0x164`[cHTReqDistEn]=1.
- `D18F0x164`[DstLnk] is the ganged link to the node that is halfway to `D18F0x164`[DstNode].
- If there are 2 ganged links where there are 2 ganged link hops to DstNode, then DstLnk must be set such that victim traffic from only 1 node must travel on a link in the same direction. See Figure 9; boxes are nodes; thick is a ganged link; thin is an unganged link. Either of the following schemes are acceptable:
 - Horizontal then vertical:
 - Node 0 -> Node 1 -> Node 2. (Node0: DstLnk=a, DstNode=2)
 - Node 1 -> Node 0 -> Node 3. (Node1: DstLnk=a, DstNode=3)
 - Node 2 -> Node 3 -> Node 0. (Node2: DstLnk=c, DstNode=0)
 - Node 3 -> Node 2 -> Node 1. (Node3: DstLnk=c, DstNode=1)
 - Vertical then horizontal:
 - Node 0 -> Node 3 -> Node 2. (Node0: DstLnk=d, DstNode=2)

- Node 1 -> Node 2 -> Node 3. (Node1: DstLnk=b, DstNode=3)
- Node 2 -> Node 1 -> Node 0. (Node2: DstLnk=b, DstNode=0)
- Node 3 -> Node 0 -> Node 1. (Node3: DstLnk=d, DstNode=1)
- For the G34r1 package, the external 8 bit link that is marked as ganged for performance reasons with a disabled internal 8 bit internal link should not be eligible for victim distribution. See 2.12.1.5 [Link Mapping Between Package and Node].




Figure 9: Example 4-Node System in Full Topology

2.9.3.2.5 Display Refresh And IFCM

Display refresh traffic is traffic generated by UMA graphics chipsets. It targets system memory for the purpose of refreshing the display. Link display refresh packets are defined as follows:

1. IO-initiated, non-posted read requests with the isochronous bit, PassPW bit, and RespPassPW bit set, and the coherent bit cleared. The SeqID must be zero and the request must be addressed outside the GART aperture.
2. The corresponding response to these requests.

The NB prioritizes these packets such that display refresh latency and bandwidth goals may be met. To support display refresh traffic, [D18F0x68 \[Link Transaction Control\]\[DispRefModeEn\]](#) is set.

Alternatively, if supported by the chipset, link-defined isochronous flow control mode (IFCM) may be employed. IFCM is enabled through [D18F0x\[E4,C4,A4,84\] \[Link Control\]\[IsocEn\]](#). If this bit is set for any link, then [D18F0x68\[DispRefModeEn\]](#) must be clear.

- The processor does not support peer-to-peer accesses in isochronous virtual channels. Upstream isochronous requests that target IO space are passed to the IO device in the base channel (the Isoc bit in the request packet is low); however, the Isoc bit in the downstream response to the requester is still set in such a case.
- In non-IFCM, the link-defined Isoc bit in the request packet is cleared as it is reflected downstream in a peer-to-peer access as well.

See [D18F0x\[F0,D0,B0,90\] \[Link Base Channel Buffer Count\]](#), [D18F3x6C \[Data Buffer Count\]](#), [D18F3x70 \[SRI to XBAR Command Buffer Count\]](#), and [D18F3x74 \[XBAR to SRI Command Buffer Count\]](#) for IFCM buffer requirements and [D18F3x140 \[SRI to XCS Token Count\]](#) for IFCM and display refresh token requirements.

If an IOMMU is present on a link, IFCM must be enabled for that link.

2.9.3.2.5.1 Recommended Buffer Count Settings Overview

When changing from the recommended settings, see the register programming requirements in the definition of each register. Some chipsets may further optimize these settings for their platform. If values other than the recommended settings are used, see the register requirements in the definition of each register. Table 23 defines commonly used terms for the following tables.

Table 23: Link Definitions

Term	Definition
Gen1	Refers to older revisions of the link specification and, in particular, link data rates from 0.4 to 2.0 GT/s. See 2.12 [Links].
Gen3	Refers to link data rates from 2.4 to 6.4GT/s, as specified by revision 3.00 of the link specification. See 2.12 [Links]. Gen3 = ($\{D18F0x[FC,DC,BC,9C][Freq[4]], D18F0x[E8,C8,A8,88][Freq[3:0]]\} \geq 7h$).
G34r1	G34r1 package. G34r1 = (CPUID Fn8000_0001_EBX[PkgType]==0011b).
C32r1	C32r1 package. C32r1 = (CPUID Fn8000_0001_EBX[PkgType]==0101b).
AM3r2	AM3r2 package. AM3r2 = (CPUID Fn8000_0001_EBX[PkgType]==0001b).
SCM	1 or 2 C32r1 sockets or 1 AM3r2 socket. SCM = (C32r1 AM3r2).
MCM1	1 G34r1 socket.
MCM2h	2 G34r1 socket half populated. See Figure 5: [Two Processor G34r1 Topology].
MCM2	2 G34r1 socket fully populated. See Figure 5: [Two Processor G34r1 Topology].
MCM4h	4 G34r1 socket half populated. See Figure 6: [Four Processor G34r1 Maximum Performance Topology] and Figure 6: [Four Processor G34r1 Maximum Performance Topology].
MCM4	4 G34r1 socket fully populated. See Figure 6: [Four Processor G34r1 Maximum Performance Topology] and Figure 6: [Four Processor G34r1 Maximum Performance Topology].
LinkConnected	LinkConnected = D18F0x[F8,D8,B8,98][LinkCon].
LinkGang	Ganged = D18F0x[17C:170][Ganged]. See D18F0x[18C:170][Ganged].
IoLink	IoLink = D18F0x[F8,D8,B8,98][NC].
NoIoLink	Indicates that none of the links are configured as an IO link.
IOMMU	Indicates the presence of an IOMMU device on the IOH. IOMMU uses the isochronous flow control channel. If an IOMMU is present, D18F0x[E4,C4,A4,84][IsocEn] must be set for all links.
IFCM	Isochronous Flow Control Mode. IFCM = D18F0x[E4,C4,A4,84][IsocEn].
NFCM	Non-Isochronous Flow Control Mode. NFCM = ~IFCM & ~D18F0x68[DispRefModeEn].
UmaDr	UMA Display Refresh. UmaDr = D18F0x68[DispRefModeEn].
UmaIfcm	UmaIfcm = ~D18F0x68[DispRefModeEn] & IFCM & D18F3x158[LnkToXcsDRToken]!=0.
NonUmaIfcm	NonUmaIfcm = ~D18F0x68[DispRefModeEn] & IFCM & ~D18F3x158[LnkToXcsDRToken].
IoUnganged	IoUnganged = IoLink & ~LinkGang.
CohGanged	CohGanged = ~IoLink & LinkGang & ~IOMMU.
CohIommu	CohIommu = ~IoLink & LinkGang & IOMMU.

Table 23: Link Definitions

Term	Definition
CohUnganged	CohUnganged = \sim IoLink & \sim LinkGang.
32Byte	32 byte, and optionally 64 byte, display refresh requests are generated by the graphics engine. \sim 32byte indicates that only 64 byte display refresh requests are generated by the graphics engine.
MultiLink	More than 1 link exists. MultiLink = (D18F0x80[CapPtr]! \neq 00h). See D18F0x[E0,C0,A0,80]. Multilink does not indicate whether more than 1 link is connected.
SingleLink	Only 1 link exists and is by definition the IO link. SingleLink = (D18F0x80[CapPtr] \neq 00h). See D18F0x[E0,C0,A0,80].
ProcCnt	Number of processors. After BIOS initializes D18F0x60[NodeCnt], ProcCnt is D18F0x60[NodeCnt]/2 for G34 MCM and D18F0x60[NodeCnt] otherwise. ProcCnt = (G34r1 ? (D18F0x60[NodeCnt]/2) : D18F0x60[NodeCnt]).
PrbFltrEn	Probe filter is enabled. PrbFltrEn = (D18F3x1D4[PFMode]! \neq 00b).
LinkFreq	Link frequency in MHz. LinkFreq = {D18F0x[FC,DC,BC,9C][Freq[4]], D18F0x[E8,C8,A8,88][Freq[3:0]]}.

2.9.4 The Level 3 Cache (L3)

The NB includes an L3 cache.

Table 24: L3 Controller Definitions

Term	Definition
L3qSize	L3 command buffer queue size. L3qSize = 32.

2.9.4.1 Probe Filter

The probe filter provides filtering of broadcast probes to improve link bandwidth and performance for multi-node systems. It uses a portion of the L3 data cache as a directory to track cache lines in the system. The probe filter directory on each node tracks addresses mapped to the local DRAM on that node. The probe filter is enabled by programming D18F3x1D4 [Probe Filter Control] according to 2.9.4.2 [Probe Filter and ATM Mode Initialization Sequence].

Probe filter configuration requirements:

- The probe filter is supported only for multiple node systems.
- The probe filter is recommended to be enabled if 2 nodes and the MEMCLK frequency for all DCTs in the system is at least 533 MHz.
- The probe filter is recommended to be enabled if \geq 3 nodes.
- The probe filter should be enabled or disabled on all nodes in the system.
- The 8-way probe filter must only be enabled if the total L3 capacity on a node (prior to enabled the probe filter) is equal to 8 MB and no subcache is less than 2 MB.

See 2.9.4.2 [Probe Filter and ATM Mode Initialization Sequence].

2.9.4.1.1 Probe Filter Errors

The L3 data cache used to store probe filter directory data is ECC protected. Single-bit (correctable) errors are corrected and logged in **D18F3x4C** [MCA NB Status High]. Multi-bit (uncorrectable) errors are logged in **D18F3x4C**, and the faulty probe filter location is no longer used. The directory is periodically scrubbed and the L3 scrubber may be redirected to scrub a location on correctable errors.

Uncorrectable ECC errors in the directory are not fatal, do not set **D18F3x4C**[UC], and do not cause machine check exceptions, but the directory entry is no longer reliable. The L3 reports uncorrectable ECC errors to the probe filter. Once such an error occurs, the faulty location index in the directory is no longer used and all requests which map to that index issue broadcast probes. **D18F3x1D4**[PFeccError] is set on an uncorrectable ECC error and can only be cleared by software. An interrupt may be generated on uncorrectable ECC errors by configuring **D18F3x1D4**[PFErrInt]. An interrupt may be generated on uncorrectable ECC errors by configuring **D18F3x1D4**[PFErrInt].

2.9.4.2 Probe Filter and ATM Mode Initialization Sequence

Perform the following steps to enable the probe filter or enable ATM mode:

1. BSC: Disable the L3 and DRAM scrubbers on all nodes in the system:
 - **D**[1F:18]F3x58[L3Scrub]=00h.
 - **D**[1F:18]F3x58[DramScrub]=00h.
 - **D**[1F:18]F3x5C[ScrubReDirEn]=0.
2. BSC: Wait 40us for outstanding scrub requests to complete.
3. All cores in the system: Allowed to be done sequentially. Disable all cache activity in the system and empty caches for all active cores in the system:
 - A. IF (**MSRC001_102B**[CombineCr0Cd]==0) THEN **MSRC001_102B**[CombineCr0Cd]=1. ENDIF.
 - B. **CR0**[CD]=1.
 - C. Issue **WBINVD**.
 - D. If (probe filter configuration needs to be enabled) then **MSRC001_102A**[ProbeFilterSupEn]=1.
4. BSC: Change configuration state for all active nodes:
 - Perform the following if the probe filter configuration needs to be enabled:
 1. The following register settings must be used in systems where the probe filter is enabled. There is no implied ordering for the following settings.
 - **D**[1F:18]F2x1B0[CohPrefPrbLmt]=000b.
 2. Set **D**[1F:18]F3x1C4[L3TagInit]=1.
 3. Wait for **D**[1F:18]F3x1C4[L3TagInit]==0.
 4. Initialize **D18F3x1D4**:
 - **D**[1F:18]F3x1D4[PFMode]: BIOS: IF (**D18F3x1C4**[15:0]==CCCCh) THEN 11b ELSE 10b ENDIF.
 - **D**[1F:18]F3x1D4[PFWayNum]: BIOS: 10b.
 - **D**[1F:18]F3x1D4[PFSubCacheSize0, PFSubCacheSize1, PFSubCacheSize2, PFSubCacheSize3]: BIOS: IF (**D18F3x1C4**[15:0]==CCCCh) THEN 01b ELSE 00b ENDIF.
 - **D**[1F:18]F3x1D4[PFSubCacheEn]: BIOS: Fh.
 - **D**[1F:18]F3x1D4[PFWayHashEn]: BIOS: 1.
 - **D**[1F:18]F3x1D4[PFPreferredSOREpl]: BIOS: IF (**D18F3x1C4**[15:0]==CCCCh) THEN 10b ELSE 00b ENDIF.
 - **D**[1F:18]F3x1D4[PFLoIndexHashEn]: BIOS: 1.
 5. Wait for **D**[1F:18]F3x1D4[PFInitDone]==1. // PF directory init is complete.
 - Perform the following if the ATM configuration needs to be enabled:

1. D[1F:18]F0x68[ATMMoDeEn]=1.
2. D[1F:18]F3x1B8[L3ATMMoDeEn]=1.
5. All cores in sequence: Enable all cache activity in the system by clearing CR0[CD] for all active cores in the system. Leave MSRC001_102B[CombineCr0Cd]=1 on all active cores.
6. BSC: Restore L3 and DRAM scrubber register values on all nodes in the system as modified by step 1.

2.9.4.3 L3 Cache Partitioning

In order to minimize the displacement of L3 data of one compute unit by another compute unit, L2 victim traffic from each compute unit can be directed to a specified set of subcaches of the L3; see D18F4x1D4 [Compute Unit Based L3 Cache Partitioning]. See D18F3x1C4 for the size of each of the 4 L3 subcaches. This mechanism will direct L2 victims according to D18F4x1D4. D18F4x1D4 can be configured while L2 victims exist and are being generated. L2 victim traffic existing or generated at the time that D18F4x1D4 is configured may not obey the updated configuration; all or nearly all L2 victim traffic existing or generated after the time that D18F4x1D4 is configured will obey the updated configuration.

Programming requirements:

- Both BAN and L3 cache partitioning are methods to influence L3 allocation policy, they interfere with each other, and thus are not recommended to be enabled simultaneously. Before L3 cache partitioning is enabled for any compute unit via D18F4x1D4, BAN mode is recommended to be disabled by (D18F3x1B8[L3BanMode]==00b).
- There are no restrictions for cache state or activity that can be in flight at the time of the configuration writes.
- The configuration of L3 cache partitioning and BAN mode is node specific and may vary by node.

2.9.5 Memory Scrubbers

The processor includes memory scrubbers specified in D18F3x58 [Scrub Rate Control], D18F3x5C, and D18F3x60. The scrubbers ensure that all cachelines in memory within or connected to the processor are periodically read and, if correctable errors are discovered, they are corrected. The system memory scrubber is also employed as specified in D18F3xB0 [On-Line Spare Control][SwapEn0].

For recommendations on scrub rates, see 2.13.1.8 [Scrub Rate Considerations].

The scrub rate is specified as the time between successive scrub events. A scrub event occurs when a line of memory is checked for errors; the amount of memory that is checked varies based on the memory block (see field descriptions).

The time required to fully scrub the memory of a node is determined as:

- Time = ((memory size in bytes)/64) * (Scrub Rate).
- E.g. If a node contains 1GB of system memory and DramScrub=5.24 ms, then all of the system memory of the node is scrubbed about once every 23 hours.

2.10 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.

BIOS reads D18F5x84[DctEn] and D18F3xE8[MultiNodeCpu] to determine the DCT to DDR channel mapping as follows:

- For single-node products, DCT0 and DCT1 respectively control channels A and B.

- For dual-node products, DCT0 and DCT1 of internal node 0 respectively control channels A and B. DCT0 and DCT1 of internal node 1 respectively control channels C and D.

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 [2.10.6 \[Memory Interleaving Modes\]](#) and [2.10.7 \[Memory Hoisting\]](#).

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

- All DIMMs connected to a node are required to operate at the same MEMCLK frequency, regardless of the channel. Both DCTs must be programmed to the same frequency.
- Mixing of unbuffered, registered, and/or load reduced DIMMs within a system is not supported.
- Mixing of ECC and non-ECC DIMMs within a system is not supported.

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 25: DCT Definitions

Term	Definition								
SR	Single Rank								
DR	Dual Rank								
QR	Quad Rank								
LR	Load Reduced DIMM								
MRS	JEDEC defined DRAM Mode Register Set.								
NP	No DIMM populated								
DataMaskMbType	<p>Motherboard type for processor Data Mask pins.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No connect</td> </tr> <tr> <td>01b</td> <td>Pins are routed per DM rules</td> </tr> <tr> <td>10b</td> <td>Pins are routed per DQS rules</td> </tr> </tbody> </table>	Bits	Description	00b	No connect	01b	Pins are routed per DM rules	10b	Pins are routed per DQS rules
Bits	Description								
00b	No connect								
01b	Pins are routed per DM rules								
10b	Pins are routed per DQS rules								
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									
DIMM2									
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								
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.								
DeviceWidth	SPDByte[7][2:0] of the DIMM being configured.								

Table 25: DCT Definitions

Term	Definition
AutoSelfRefresh	SPDByte[31][2] of the DIMM being configured.
ExtendedTemperature-Range	SPDByte[31][0] of the DIMM being configured.
MinimumModuleDelay	SPDByte[94, 92, 90][6:0], indexed by VDDIO, of the LRDIMM being configured.
MaximumModuleDelay	SPDByte[95, 93, 91][6:0], indexed by VDDIO, of the LRDIMM being configured.
DdrRate	The DDR data rate (MT/s).
DdrVDDIO	DDR VDDIO in V.
SODIMM	SPDByte[3][3:0] == 3h. DCT is configured for SODIMM if (D18F2x90_dct[1:0][UnbuffDimm]==1) and all enabled ranks have (D18F2x[6C:60]_dct[1:0][RankDef]==0).
UDIMM	SPDByte[3][3:0] == 2h. DCT is configured for UDIMM if (D18F2x90_dct[1:0][UnbuffDimm]==1) and all enabled ranks have (D18F2x[6C:60]_dct[1:0][RankDef]==0).
RDIMM	SPDByte[3][3:0] == 1h. DCT is configured for RDIMM if (D18F2x90_dct[1:0][UnbuffDimm]==0) and all enabled ranks have (D18F2x[6C:60]_dct[1:0][RankDef]==0).
LRDIMM	SPDByte[3][3:0] == Bh. DCT is configured for LRDIMM if (D18F2x90_dct[1:0][UnbuffDimm]==0) and any enabled rank has (D18F2x[6C:60]_dct[1:0][RankDef]!=0).

Table 26: DDR3 UDIMM Maximum Frequency Support for G34

DIMM Slots/Ch	DIMMs	DIMMs		Frequency ¹ (MT/s)		
		SR	DR	1.5V	1.35V	1.25V
1	1	1	-	1866	1600	1333
		-	1	1866	1600	1333
2	1	1	-	1600	1333	1333
		-	1	1600	1333	1333
	2	2	-	1600	1333	1333
		1	1	1333	1333	1066
3	1	1	-	1600	1333	1333
		-	1	1600	1333	1333
	2	2	-	1333	1333	1333
		1	1	1333	1333	1066
		-	2	1333	1333	1066

1. Population restrictions (including the order for partially populated channels) may apply. See Table 35.

Table 27: DDR3 RDIMM Maximum Frequency Support for G34

DIMM Slots/Ch	DIMMs	DIMMs			Frequency ¹ (MT/s)		
		SR	DR	QR	1.5V	1.35V	1.25V
1	1	1	-	-	1866	1600	1333
		-	1	-	1866	1600	1333
		-	-	1	1333	1333	1066
2	1	1	-	-	1600	1333	1333
		-	1	-	1600	1333	1333
		-	-	1	1333	1066	1066
	2	2	-	-	1600	1333	1333
		1	1	-	1600	1333	1066
		1	-	1	1066	800	800
		-	2	-	1600	1333	1066
		-	1	1	1066	800	800
		-	-	2	1066	1066	800
3	1	1	-	-	1600	1333	1333
		-	1	-	1600	1333	1333
		-	-	1	1066	1066	800
	2	2	-	-	1600	1333	1333
		1	1	-	1600	1333	1066
		1	-	1	800	800	800
		-	2	-	1600	1333	1066
		-	1	1	800	800	800
		-	-	2	1066	1066	800
	3	3	-	-	1066	1066	800
		2	1	-	1066	800	667
		2	-	1	800	800	667
		1	2	-	1066	800	667
		1	1	1	800	800	667
		-	3	-	1066	800	667
-		2	1	800	800	667	

1. Population restrictions (including the order for partially populated channels) may apply. See [Table 36](#).

Table 28: DDR3 LRDIMM Maximum Frequency Support for G34

DIMM Slots/Ch	DIMMs	DIMMs	Frequency ¹ (MT/s)		
		LR	1.5V	1.35V	1.25V
1	1	1	-	1333	1333
2	1	1	-	1333	1333
	2	2	-	1333	1066

Table 28: DDR3 LRDIMM Maximum Frequency Support for G34

DIMM Slots/Ch	DIMMs	DIMMs	Frequency ¹ (MT/s)		
		LR	1.5V	1.35V	1.25V
3	1	1	-	1333	1333
	2	2	-	1333	1066
	3	3	N/A-1066 ³	800 ² -1066 ³	667 ²

- Population restrictions (including the order for partially populated channels) may apply. See [Table 37](#).
- The indicated frequencies represent the current plan; however, there is a possibility that these frequencies may improve by one speed grade.
- Support varies by vendor. See your AMD representative for more information.

Table 29: DDR3 UDIMM Maximum Frequency Support for C32

DIMM Slots/Ch	DIMMs	DIMMs		Frequency ¹ (MT/s)		
		SR	DR	1.5V	1.35V	1.25V
1	1	1	-	1600 ²	1333 ²	1333
		-	1	1600 ²	1333 ²	1333
2	1	1	-	1600	1333	1333
		-	1	1600	1333	1333
	2	2	-	1333	1333	1066 ²
		1	1	1066 ²	1066 ²	1066
3	1	1	-	1600	1333	1333
		-	1	1600	1333	1333
	2	2	-	1333	1333	1066 ²
		1	1	1066 ²	1066 ²	1066
		-	2	1066 ²	1066 ²	1066

- Population restrictions (including the order for partially populated channels) may apply. See [Table 35](#).
- The indicated frequencies represent the current plan; however, there is a possibility that these frequencies may improve by one speed grade.

Table 30: DDR3 RDIMM Maximum Frequency Support for C32

DIMM Slots/Ch	DIMMs	DIMMs			Frequency ¹ (MT/s)		
		SR	DR	QR	1.5V	1.35V	1.25V
1	1	1	-	-	1600 ²	1333 ²	1333
		-	1	-	1600 ²	1333 ²	1333
		-	-	1	1333	1066 ²	1066

Table 30: DDR3 RDIMM Maximum Frequency Support for C32

DIMM Slots/Ch	DIMMs	DIMMs			Frequency ¹ (MT/s)		
		SR	DR	QR	1.5V	1.35V	1.25V
2	1	1	-	-	1600	1333	1333
		-	1	-	1600	1333	1333
		-	-	1	1333	1066 ²	1066
	2	2	-	-	1333	1333	1066 ²
		1	1	-	1066 ²	1066 ²	1066
		1	-	1	800 ²	800	800
		-	2	-	1066 ²	1066 ²	1066
		-	1	1	800 ²	800	800
		-	-	2	800 ²	800	800
3	1	1	-	-	1600	1333	1333
		-	1	-	1600	1333	1333
		-	-	1	1066	800 ²	800
	2	2	-	-	1333	1333	1066 ²
		1	1	-	1066 ²	1066 ²	1066
		1	-	1	800 ²	800	800
		-	2	-	1066 ²	1066 ²	1066
		-	1	1	800 ²	800	800
		-	-	2	800 ²	800	800
	3	3	-	-	1066 ²	800 ²	800
		2	1	-	800 ²	800	667 ²
		2	-	1	667 ²	667 ²	667
		1	2	-	800 ²	800	667 ²
		1	1	1	667 ²	667 ²	667
		-	3	-	800 ²	800	667 ²
-		2	1	667 ²	667 ²	667	

1. Population restrictions (including the order for partially populated channels) may apply. See [Table 36](#).
2. The indicated frequencies represent the current plan; however, there is a possibility that these frequencies may improve by one speed grade.

Table 31: DDR3 LRDIMM Maximum Frequency Support for C32

DIMM Slots/Ch	DIMMs	DIMMs	Frequency ¹ (MT/s)		
		LR	1.5V	1.35V	1.25V
1	1	1	-	1333 ²	1333
2	1	1	-	1333	1333
	2	2	-	1066 ²	1066

Table 31: DDR3 LRDIMM Maximum Frequency Support for C32

DIMM Slots/Ch	DIMMs	DIMMs	Frequency ¹ (MT/s)		
		LR	1.5V	1.35V	1.25V
3	1	1	-	1333	1333
	2	2	-	1066 ²	1066
	3	3	-	800	667 ²

1. Population restrictions (including the order for partially populated channels) may apply. See [Table 37](#).

2. The indicated frequencies represent the current plan; however, there is a possibility that these frequencies may improve by one speed grade.

Table 32: DDR3 UDIMM Maximum Frequency Support for AM3r2 Desktop

DIMM Slots/Ch	DIMMs	DIMMs		Frequency ¹ (MT/s)
		SR	DR	1.5V
1	1	1	-	1866
		-	1	1866
2	1	1	-	1600 ²
		-	1	1600 ²
	2	2	-	1600
		1	1	1333 ²
-	2	1333 ²		

1. Population restrictions (including the order for partially populated channels) may apply. See [Table 38](#).

2. The indicated frequencies represent the current plan; however, there is a possibility that these frequencies may improve by one speed grade.

Table 33: DDR3 UDIMM Maximum Frequency Support for AM3r2 Microserver

DIMM Slots/Ch	DIMMs	DIMMs		Frequency ¹ (MT/s)	
		SR	DR	1.5V	1.35V
1	1	1	-	1866	1333
		-	1	1866	1333
2	1	1	-	1600	1333
		-	1	1600	1333
	2	2	-	1600	1333
		1	1	1333	1066
-	2	1333	1066		

1. Population restrictions (including the order for partially populated channels) may apply. See [Table 38](#).

2.

Table 34: DDR3 SODIMM Maximum Frequency Support for C32/AM3r2 Microserver

DIMM Slots/Ch	DIMMs	DIMMs		Frequency ¹ (MT/s)	
		SR	DR	1.5V	1.35V
1	1	1	-	1600	1333
		-	1	1600	1333
2	1	1	-	1333	1333
		-	1	1333	1333
	2	2	-	1333	1333
		1	1	1333	1066
		-	2	1333	1066

1. Population restrictions (including the order for partially populated channels) may apply. See [Table 38](#).

2.

The tables below list the DIMM populations as supported by the processor. DIMMs must be populated from farthest slot to closest slot to the processor on a per channel basis when a daisy chain topology is used. DIMMs are populated on the outer slots first on a 3 DIMM/Ch tee or star topology.

Table 35: DDR3 UDIMM Population Support G34/C32

DIMM Slots/Ch	DIMM0	DIMM1	DIMM2
1	SR/DR	N/A	N/A
2	-	SR/DR	N/A
	SR/DR	SR/DR	N/A
3	-	-	SR/DR
	SR/DR	-	SR/DR

Table 36: DDR3 RDIMM Population Support G34/C32

DIMM Slots/Ch	DIMM0	DIMM1	DIMM2
1	SR/DR/QR	N/A	N/A
2	-	SR/DR/QR	N/A
	SR/DR/QR	SR/DR/QR	N/A
3	-	-	SR/DR
	-	QR	-
	-	QR	SR/DR
	SR/DR	-	SR/DR
	SR/DR	SR/DR/QR	SR/DR

Table 37: DDR3 LRDIMM Population Support G34/C32

DIMM Slots/Ch	DIMM0	DIMM1	DIMM2
1	LR	N/A	N/A
2	-	LR	N/A
	LR	LR	N/A
3	-	-	LR
	LR	-	LR
	LR	LR	LR

Table 38: DDR3 UDIMM Population Support AM3r2

DIMM Slots/Ch	DIMM0	DIMM1
1	SR/DR	N/A
2	-	SR/DR
	SR/DR	SR/DR

Table 39: DDR3 SODIMM Population Support C32/AM3r2

DIMM Slots/Ch	DIMM0	DIMM1
1	SR/DR	N/A
2	-	SR/DR
	SR/DR	SR/DR

2.10.1 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.

2.10.2 DDR Pad to Processor Pin Mapping

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

Table 40: Package pin mapping

Pad	Pin			
	G34r1 Internal Node 0	G34r1 Internal Node 1	C32r1	AM3r2
MEMCLK0_H[0]	MA_CLK_H[0]	MC_CLK_H[0]	MA_CLK_H[0]	MA_CLK_H[2]
MEMCLK0_H[1]	MA_CLK_H[1]	MC_CLK_H[1]	MA_CLK_H[1]	MA_CLK_H[4]
MEMCLK0_H[2]	MA_CLK_H[2]	MC_CLK_H[2]	MA_CLK_H[2]	MA_CLK_H[5]
MEMCLK0_H[3]	MA_CLK_H[3]	MC_CLK_H[3]	MA_CLK_H[3]	MA_CLK_H[3]
MEMCLK0_H[4]	MA_CLK_H[4]	MC_CLK_H[4]	-	-
MEMCLK0_H[5]	-	-	-	-
MEMCLK1_H[0]	MB_CLK_H[0]	MD_CLK_H[0]	MB_CLK_H[0]	MB_CLK_H[2]
MEMCLK1_H[1]	MB_CLK_H[1]	MD_CLK_H[1]	MB_CLK_H[1]	MB_CLK_H[4]
MEMCLK1_H[2]	MB_CLK_H[2]	MD_CLK_H[2]	MB_CLK_H[2]	MB_CLK_H[5]
MEMCLK1_H[3]	MB_CLK_H[3]	MD_CLK_H[3]	MB_CLK_H[3]	MB_CLK_H[3]
MEMCLK1_H[4]	MB_CLK_H[4]	MD_CLK_H[4]	-	-
MEMCLK1_H[5]	-	-	-	-
MEMCS0_L[0]	MA0_CS_L[0]	MC0_CS_L[0]	MA0_CS_L[0]	MA0_CS_L[0]
MEMCS0_L[1]	MA0_CS_L[1]	MC0_CS_L[1]	MA0_CS_L[1]	MA0_CS_L[1]
MEMCS0_L[2]	MA1_CS_L[0]	MC1_CS_L[0]	MA1_CS_L[0]	MA1_CS_L[0]
MEMCS0_L[3]	MA1_CS_L[1]	MC1_CS_L[1]	MA1_CS_L[1]	MA1_CS_L[1]
MEMCS0_L[4]	MA2_CS_L[0]	MC2_CS_L[0]	MA0_CS_L[2]	-
MEMCS0_L[5]	MA2_CS_L[1]	MC2_CS_L[1]	MA0_CS_L[3]	-
MEMCS0_L[6]	MA3_CS_L[0]	MC3_CS_L[0]	MA1_CS_L[2]	-
MEMCS0_L[7]	MA3_CS_L[1]	MC3_CS_L[1]	MA1_CS_L[3]	-
MEMCS1_L[0]	MB0_CS_L[0]	MD0_CS_L[0]	MB0_CS_L[0]	MB0_CS_L[0]
MEMCS1_L[1]	MB0_CS_L[1]	MD0_CS_L[1]	MB0_CS_L[1]	MB0_CS_L[1]
MEMCS1_L[2]	MB1_CS_L[0]	MD1_CS_L[0]	MB1_CS_L[0]	MB1_CS_L[0]
MEMCS1_L[3]	MB1_CS_L[1]	MD1_CS_L[1]	MB1_CS_L[1]	MB1_CS_L[1]
MEMCS1_L[4]	MB2_CS_L[0]	MD2_CS_L[0]	MB0_CS_L[2]	-
MEMCS1_L[5]	MB2_CS_L[1]	MD2_CS_L[1]	MB0_CS_L[3]	-
MEMCS1_L[6]	MB3_CS_L[0]	MD3_CS_L[0]	MB1_CS_L[2]	-
MEMCS1_L[7]	MB3_CS_L[1]	MD3_CS_L[1]	MB1_CS_L[3]	-
MEMODT0[0]	MA0_ODT[0]	MC0_ODT[0]	MA0_ODT[0]	MA0_ODT[0]
MEMODT0[1]	MA1_ODT[0]	MC1_ODT[0]	MA1_ODT[0]	MA1_ODT[0]
MEMODT0[2]	MA2_ODT[0]	MC2_ODT[0]	MA0_ODT[1]	MA0_ODT[1]
MEMODT0[3]	MA3_ODT[0]	MC3_ODT[0]	MA1_ODT[1]	MA1_ODT[1]
MEMODT1[0]	MB0_ODT[0]	MD0_ODT[0]	MB0_ODT[0]	MB0_ODT[0]
MEMODT1[1]	MB1_ODT[0]	MD1_ODT[0]	MB1_ODT[0]	MB1_ODT[0]
MEMODT1[2]	MB2_ODT[0]	MD2_ODT[0]	MB0_ODT[1]	MB0_ODT[1]
MEMODT1[3]	MB3_ODT[0]	MD3_ODT[0]	MB1_ODT[1]	MB1_ODT[1]
MEMCKE0[0]	MA_CKE[0]	MC_CKE[0]	MA_CKE[0]	MA_CKE[0]
MEMCKE0[1]	MA_CKE[1]	MC_CKE[1]	MA_CKE[1]	MA_CKE[1]
MEMCKE1[0]	MB_CKE[0]	MD_CKE[0]	MB_CKE[0]	MB_CKE[0]
MEMCKE1[1]	MB_CKE[1]	MD_CKE[1]	MB_CKE[1]	MB_CKE[1]

1. For differential pins, only positive polarity pins are shown; negative polarity pins have corresponding mapping and are controlled by the same pad.

2.10.3 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][EnDr-amInit]` or `D18F2x90_dct[1:0][InitDram]` is set to 1.

See 2.10.5.7 [DRAM Device and Controller Initialization] and 2.3.3 [Using L2 Cache as General Storage During Boot].

2.10.4 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.

2.10.5 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 2.10.5.9 [DRAM Channel Disable].

1. Configure the DDR supply voltage regulator. See 2.10.5.1.
2. Force NB P-state to NBP0. See 2.10.5.2.
3. DDR phy initialization. See 2.10.5.3.
4. 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 2.10.5.4.
 - b. Program Non-SPD configuration. See 2.10.5.5.
 - c. Program DCT training specific configuration. See 2.10.5.6.
 - d. Program the remaining DCT registers not covered by an explicit sequence dependency.
 - e. DRAM device initialization. See 2.10.5.7.
 - 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 2.10.5.4 and 2.10.5.5 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 steps 5 and 6 below) from non-volatile storage.
 - d. Continue at step 8.
5. DRAM write levelization training. See 2.10.5.8.1.
6. DRAM data training.
 - A. DQS receiver enable training.
 - a. DQS receiver enable training. See 2.10.5.8.2
 - b. Program `D18F2x9C_x0D0F_E003_dct[1:0][DisAutoComp, DisablePredriverCal] = {0b, 1b}`.
 - c. DQS receiver enable cycle training. See 2.10.5.8.3.

- B. DQS position training. See 2.10.5.8.4.
- C. MaxRdLatency training. See 2.10.5.8.5.1.
- 7. NB P-state specific training. For each NB P-state from NBP1 to D18F5x170[NbPstateMaxVal]:
 - A. Force the NB P-state. See 2.10.5.2.
 - B. MaxRdLatency training. See 2.10.5.8.5.1.
- 8. Release NB P-state force. See 2.10.5.2.
- 9. Program DCT for normal operation. See 2.10.5.6.
- 10. Program DRAM phy for power savings. See 2.10.5.10.

The DRAM subsystem is ready for use.

2.10.5.1 Low Voltage DDR3

JEDEC defined DDR3L and DDR3U devices are supported. Platforms that support 1.35V or 1.25V operation should power on VDDIO at 1.35V until operating voltage is determined by reading the SPD ROM of all DIMMs. BIOS should not operate DIMMs at voltages higher than supported as indicated by SPD.

The recommended BIOS configuration sequence is as follows:

1. BIOS reads the SPD ROM of all DIMMs to determine the common operating voltage.
2. BIOS configures VDDIO to match the lowest common supported voltage based on the SPD values. See platform specific documentation for changing the voltage.
3. Additional derating of the DDR speed may be necessary for reliable operation at lower voltage. See 2.10 for information on voltage and electrical load specific maximum speed.

2.10.5.2 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 force the processor to the desired NB P-state:

1. Program the NB P-state specific configuration registers for the target NB P-state.
 - D18F2x210_dct[1:0]_nbp[3:0][MaxRdLatency, DataTxFifoWrDly, RdPtrInit].
2. Program D18F5x170 to transition the NB P-state:
 - SwNbPstateLoDis = NbPstateDisOnP0 = NbPstateThreshold = 0.
 - NbPstateLo = NbPstateMaxVal.
3. Wait for D18F5x174[CurNbPstate] to equal NbPstateLo.
4. Program D18F5x170 to force the NB P-state:
 - NbPstateHi = target NB P-state.
 - SwNbPstateLoDis = 1
5. Wait for D18F5x174[CurNbPstate] to equal the target NB P-state.

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

6. Restore the initial D18F5x170[SwNbPstateLoDis, NbPstateDisOnP0] values.
7. Restore the initial D18F5x170[NbPstateThreshold, NbPstateHi] values.

See also 2.5.2.2 [NB P-states].

2.10.5.3 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_x0000_000B_dct[1:0]` = 80000000h.
2. Program `D18F2x9C_x0D0F_E013_dct[1:0]` = 0118h.
3. Phy Voltage Level Programming. See 2.10.5.3.1.
4. DRAM channel frequency change. See 2.10.5.3.2.
5. Phy fence programming. See 2.10.5.3.3.
6. Phy compensation initialization. See 2.10.5.3.4.

2.10.5.3.1 Phy Voltage Level Programming

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

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

See 2.10.5.1 [Low Voltage DDR3].

2.10.5.3.2 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:

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]`
 - `D18F2x9C_x0D0F_0[F,8:0]13_dct[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 2.10.5.2.

For each DCT:

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

For each DCT:

6. Program `D18F2x9C_x0D0F_E006_dct[1:0][PIILockTime]` = 0Fh.

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.

2.10.5.3.2.1 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 [2.10.5.8.1 \[Write Levelization Training\]](#) and [2.10.5.8.2 \[DQS Receiver Enable Training\]](#).

1. Force the NB P-state. See [2.10.5.2](#).
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 [2.10.5.3.2](#).
4. Phy fence programming. See [2.10.5.3.3](#).
5. Phy compensation initialization. See [2.10.5.3.4](#).
6. Program SPD configuration. See [2.10.5.4](#).
7. Program Non-SPD configuration. See [2.10.5.5](#).
8. Exit self-refresh:
 - A. Program `D18F2x90_dct[1:0][ExitSelfRef] = 1`.
 - B. Wait for `D18F2x90_dct[1:0][ExitSelfRef] = 0`.
 - C. IF (C32r1 || G34r1) THEN Program `D18F2x90_dct[1:0][DisDllShutDownSR] = 0`. ENDIF.
9. Re-program devices with frequency dependent mode register register field values. See [2.10.5.7](#).

2.10.5.3.3 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_x0000_0008_dct[1:0][FenceTrSel]=10b`.
2. Program `D18F2x9C_x0000_00[52:50]_dct[1:0]=1313_1313h`.
3. Perform phy fence training. See [2.10.5.3.3.1 \[Phy Fence Training\]](#).
4. Write the calculated fence value to `D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxDll]`.
5. Program `D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0][AlwaysEnDllClks]=001b`.
6. Program `D18F2x9C_x0000_0008_dct[1:0][FenceTrSel]=01b`.
7. Program `D18F2x9C_x0000_00[52:50]_dct[1:0]=1313_1313h`.
8. Perform phy fence training. See [2.10.5.3.3.1 \[Phy Fence Training\]](#).
9. Write the calculated fence value to `D18F2x9C_x0000_000C_dct[1:0][FenceThresholdRxDll]`.
10. Program `D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0][AlwaysEnDllClks]=000b`.
11. Program `D18F2x9C_x0000_0008_dct[1:0][FenceTrSel]=11b`.
12. Program `D18F2x9C_x0000_00[52:50]_dct[1:0]=1313_1313h`.
13. Perform phy fence training. See [2.10.5.3.3.1 \[Phy Fence Training\]](#).
14. Write the calculated fence value to `D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxPad]`.
15. IF (`D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxPad] < 16`) THEN
 - Program `D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0] = {001h,`
`D18F2x9C_x0000_000C_dct[1:0][19:16]}`
 - ELSE
 - Program `D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0] = 0000h`
 - ENDIF.

16. 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,8:0]31_dct[1:0] = {0b, Fence2_RxDll[4:0], Fence2_TxDll[4:0], Fence2_TxPad[4:0]}.
17. If motherboard routing requires CS[7:6] to adopt address timings, BIOS performs the following:
- A. Program D18F2xA8_dct[1:0][CSTimingMux67] = 1.
 - B. Program D18F2x9C_x0D0F_8021_dct[1:0]:
 - DiffTimingEn = 1.
 - IF (D18F2x9C_x0000_0004_dct[1:0][AddrCmdFineDelay] >= D18F2x9C_x0D0F_E008_dct[1:0][FenceValue]) THEN Fence = 1 ELSE Fence = 0.
 - Delay = D18F2x9C_x0000_0004_dct[1:0][AddrCmdFineDelay].
18. Reprogram D18F2x9C_x0000_0004_dct[1:0].

When resuming from S3, it is recommended that BIOS reprogram D18F2x9C_x0000_000C_dct[1:0][FenceThresholdTxDll, FenceThresholdRxDll, FenceThresholdTxPad], D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0], D18F2x9C_x0D0F_0[F,8:0]31_dct[1:0], and D18F2x9C_x0D0F_8021_dct[1:0] from values stored in from non-volatile storage instead of training.

2.10.5.3.3.1 Phy Fence Training

1. Program D18F2x9C_x0000_0008_dct[1:0][PhyFenceTrEn]=1.
2. Wait 2000 MEMCLKs.
3. Program D18F2x9C_x0000_0008_dct[1:0][PhyFenceTrEn]=0.
4. Read the phase recovery engine registers D18F2x9C_x0000_00[52:50]_dct[1:0].
5. Calculate the average of the fine delay values of all byte lanes and subtract 6. If the result is negative then the fence value is zero.

2.10.5.3.4 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_x0D0F_E003_dct[1:0][DisAutoComp, DisablePredriverCal]= { 1b, 1b}.
2. Program TxPreP/TxPreN for Data and DQS according to Table 41, Table 42, and Table 43.

- A. Program `D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0]={0000b, TxPreP, TxPreN}`.
- B. Program `D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0]={1000b, TxPreP, TxPreN}`.
- 3. Program TxPreP/TxPreN for Cmd/Addr according to [Table 50](#), [Table 51](#), and [Table 52](#).
 - 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}`.
- 4. Program TxPreP/TxPreN for Clock according to [Table 53](#), [Table 54](#), and [Table 55](#).
 - A. Program `D18F2x9C_x0D0F_2[2:0]02_dct[1:0]={1000b, TxPreP, TxPreN}`.

IF ((G34r1 || AM3r2) && !LRDIMM) THEN

Table 41: Phy Predriver Calibration Codes for Data/DQS at 1.5V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 800	000b	B6Dh
	001b	6DBh
	010b	492h
	011b	492h
1066 - 1333	000b	FFFh
	001b	924h
	010b	6DBh
	011b	6DBh
1600 - 1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	B6Dh

1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN
 See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren]
 ELSE
 See
 D18F2x9C_x0000_0000_dct[1:0][DataDrvStren]
 ENDIF.

2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and
 D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].

Table 42: Phy Predriver Calibration Codes for Data/DQS at 1.35V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 800	000b	FFFh
	001b	924h
	010b	6DBh
	011b	492h
1066 - 1333	000b	FFFh
	001b	DB6h
	010b	B6Dh
	011b	6DBh

Table 42: Phy Predriver Calibration Codes for Data/DQS at 1.35V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
1600 - 1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	DB6h

1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN
 See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren]
 ELSE
 See
 D18F2x9C_x0000_0000_dct[1:0][DataDrvStren]
 ENDIF.

2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and
 D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].

Table 43: Phy Predriver Calibration Codes for Data/DQS at 1.25V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 800	000b	FFFh
	001b	B6Dh
	010b	924h
	011b	6DBh
1066 - 1333	000b	FFFh
	001b	FFFh
	010b	DB6h
	011b	924h
1600 - 1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	FFFh

1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN
 See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren]
 ELSE
 See
 D18F2x9C_x0000_0000_dct[1:0][DataDrvStren]
 ENDIF.

2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and
 D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].

ELSE IF (C32r1 && !LRDIMM) THEN

Table 44: Phy Predriver Calibration Codes for Data/DQS at 1.5V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 800	000b	B6Dh
	001b	6DBh
	010b	492h
	011b	492h
1066	000b	FFFh
	001b	924h
	010b	6DBh
	011b	6DBh
1333	000b	B6Dh
	001b	6DBh
	010b	492h
	011b	492h
1600 - 1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	B6Dh
<p>1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0][DataDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].</p>		

Table 45: Phy Predriver Calibration Codes for Data/DQS at 1.35V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 800	000b	FFFh
	001b	924h
	010b	6DBh
	011b	492h
1066	000b	FFFh
	001b	DB6h
	010b	B6Dh
	011b	6DBh
1333	000b	FFFh
	001b	924h
	010b	6DBh
	011b	492h

Table 45: Phy Predriver Calibration Codes for Data/DQS at 1.35V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
1600 - 1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	DB6h
<p>1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0][DataDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].</p>		

Table 46: Phy Predriver Calibration Codes for Data/DQS at 1.25V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 800	000b	FFFh
	001b	B6Dh
	010b	924h
	011b	6DBh
1066	000b	FFFh
	001b	FFFh
	010b	DB6h
	011b	924h
1333	000b	B6Dh
	001b	6DBh
	010b	492h
	011b	492h
1600 - 1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	FFFh
<p>1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0][DataDrvStren] ENDIF.</p> <p>2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].</p>		

ELSE IF (LRDIMM) THEN

Table 47: Phy Predriver Calibration Codes for Data/DQS at 1.5V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 1333	000b	B6Dh
	001b	6DBh
	010b	492h
	011b	492h
1600	000b	FFFh
	001b	924h
	010b	6DBh
	011b	6DBh
1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	B6Dh

1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN
 See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren]
 ELSE
 See
 D18F2x9C_x0000_0000_dct[1:0][DataDrvStren]
 ENDIF.

2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and
 D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].

Table 48: Phy Predriver Calibration Codes for Data/DQS at 1.35V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 1333	000b	FFFh
	001b	924h
	010b	6DBh
	011b	492h
1600	000b	FFFh
	001b	DB6h
	010b	B6Dh
	011b	6DBh

Table 48: Phy Predriver Calibration Codes for Data/DQS at 1.35V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	DB6h

1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN
 See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren]
 ELSE
 See
 D18F2x9C_x0000_0000_dct[1:0][DataDrvStren]
 ENDIF.

2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and
 D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].

Table 49: Phy Predriver Calibration Codes for Data/DQS at 1.25V

DDR Rate	Drive Strength ¹	{TxPreP, TxPreN} ²
667 - 1333	000b	B6Dh
	001b	6DBh
	010b	492h
	011b	492h
1600	000b	FFFh
	001b	924h
	010b	6BDh
	011b	6BDh
1866	000b	FFFh
	001b	FFFh
	010b	FFFh
	011b	B6Dh

1. IF (D18F2x9C_x0D0F_0[F,8:0]06) THEN
 See D18F2x9C_x0000_0000_dct[1:0][DqsDrvStren]
 ELSE
 See
 D18F2x9C_x0000_0000_dct[1:0][DataDrvStren]
 ENDIF.

2. See D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] and
 D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0].

ENDIF.

Table 50: 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][CkeDrvStren] ELSEIF (D18F2x9C_x0D0F_800[A,6,2])THEN See D18F2x9C_x0000_0000_dct[1:0][CsOdtDrvStren] ELSE See D18F2x9C_x0000_0000_dct[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 51: 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

Table 51: Phy Predriver Calibration Codes for Cmd/Addr at 1.35V

DDR Rate	Drive Strength ¹	TxPreP ²	TxPreN ²
1600 - 1866	000b	101_101b	101_101b
	001b	101_101b	101_101b
	010b	100_100b	100_100b
	011b	100_100b	100_100b
1. IF (D18F2x9C_x0D0F_C002)THEN See D18F2x9C_x0000_0000_dct[1:0][CkeDrvStren] ELSEIF (D18F2x9C_x0D0F_800[A,6,2])THEN See D18F2x9C_x0000_0000_dct[1:0][CsOdtDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0][AddrCmdDrvStren] ENDIF. 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].			

Table 52: 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
1600 - 1866	000b	111_111b	110_110b
	001b	110_110b	101_101b
	010b	101_101b	100_100b
	011b	101_101b	100_100b
1. IF (D18F2x9C_x0D0F_C002)THEN See D18F2x9C_x0000_0000_dct[1:0][CkeDrvStren] ELSEIF (D18F2x9C_x0D0F_800[A,6,2])THEN See D18F2x9C_x0000_0000_dct[1:0][CsOdtDrvStren] ELSE See D18F2x9C_x0000_0000_dct[1:0][AddrCmdDrvStren] ENDIF. 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].			

Table 53: 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

1. See [D18F2x9C_x0000_0000_dct\[1:0\]\[ClkDrvStren\]](#).
2. See [D18F2x9C_x0D0F_2\[2:0\]02_dct\[1:0\]](#).

Table 54: 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
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\[1:0\]\[ClkDrvStren\]](#).
2. See [D18F2x9C_x0D0F_2\[2:0\]02_dct\[1:0\]](#).

Table 55: 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][ClkDrvStren].
2. See D18F2x9C_x0D0F_2[2:0]02_dct[1:0].

2.10.5.4 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.

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:

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

Optimal cycle time is specified for each DIMM and is used to limit or determine bus frequency.

Additionally, the SPD ROM provides values for several DIMM parameters that must be programmed by MRS or RCW commands. See 2.10.5.7 [DRAM Device and Controller Initialization]. These parameters are:

- RDIMM
 - RC3, RC4, RC5
- LRDIMM
 - F0RC2, F0RC3, F0RC4, F0RC5
 - F1RC8, F1RC11, F1RC12, F1RC13, F1RC14, F1RC15
 - F3RC8, F3RC9
 - F[10:3]RC10, F[10:3]RC11
 - F[15:14]RC[15:8,6:0]
 - MR1_dct[1:0], MR2_dct[1:0]

2.10.5.4.1 LRDIMM Module Delay Time SPD Bytes

The SPD ROM provides values for the module delay. BIOS uses this information to configure D18F2x240_dct[1:0][RdOdtTrnOnDly, RdOdtOnDuration] as follows:

- $\text{BufDatDelay} = \text{FLOOR}(\frac{(\text{SmallestModuleDelay} - \text{SynchDelay}) * (\text{MemClkFreq} / 400 \text{ MHz}) + \text{SynchDelay}}{0x40})$.
 - $\text{SmallestModuleDelay} = \text{MinimumModuleDelay} * .000125 \text{ us} * 400 \text{ MHz} * 0x40$.
 - **MinimumModuleDelay** is the minimum SPD module delay across all DIMMs on a channel.
 - $\text{SynchDelay} = (\text{F0RC2}[\text{AddrCmdPrelaunch}] ? 0x30 - (2 * \text{F1RC12}[\text{QCAPrelaunchDelay}]) : 0x20) + 0x10$.
 - SyncDelay is calculated from the SPD values of the **MinimumModuleDelay** DIMM.
- Program D18F2x240_dct[1:0][RdOdtTrnOnDly] = MAX(0, D18F2x200_dct[1:0][Tc1] - D18F2x20C_dct[1:0][Tcw1]) + BufDatDelay.
- $\text{BufDatDelaySkew} = \text{ROUND}(\frac{(\text{LargestModuleDelay} - \text{SynchDelay}) * \text{MemClkFreq} / 400 \text{ MHz} + \text{SynchDelay} - (\text{BufDatDelay} * 0x40)}{0x40})$.
 - $\text{LargestModuleDelay} = \text{MaximumModuleDelay} * .000125 \text{ us} * 400 \text{ MHz} * 0x40$.
 - **MaximumModuleDelay** is the maximum SPD module delay across all DIMMs on a channel.
 - $\text{SynchDelay} = (\text{F0RC2}[\text{AddrCmdPrelaunch}] ? 0x30 - (2 * \text{F1RC12}[\text{QCAPrelaunchDelay}]) : 0x20) + 0x10$.
 - SyncDelay is calculated from the SPD values of the **MaximumModuleDelay** DIMM.
- Program D18F2x240_dct[1:0][RdOdtOnDuration] = 6 + BufDatDelaySkew.

2.10.5.5 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 2.10.5.8 [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) = $D18F2x200_dct[1:0][Tcl] - D18F2x20C_dct[1:0][Tcwl]$.
- Read ODT Delay (ROD) = $MAX(0, D18F2x240_dct[1:0][RdOdtOnDuration] - 6)$.
- Write ODT Delay (WOD) = $MAX(0, D18F2x240_dct[1:0][WrOdtOnDuration] - 6)$.
- $WrEarly = ABS(D18F2xA8_dct[1:0][WrDqDqsEarly])/2$.
-
- $BufDatDelay = IF\ LRDIMM\ See\ 2.10.5.4.1\ [LRDIMM\ Module\ Delay\ Time\ SPD\ Bytes]\ ELSE\ 0\ ENDIF$.

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

The optimal values for $D18F2x218_dct[1:0][TrdrdSdSc, TrdrdSdDc, TrdrdDd]$ 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) = $MAX(TrdrdSdSc, 3 + (IF\ (D18F2x94_dct[1:0][FourRankRDimm1] | D18F2x94_dct[1:0][FourRankRDimm0])\ THEN\ (CEIL(CDD_{TrdrdSdDc} / 2) + 0.5)\ ELSE\ 0\ ENDIF))$.
- $TrdrdDd$ (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 $D18F2x9C_x0000_00[2A:10]_dct[1:0][DqsRcvEnGrossDelay]$ minus $D18F2x9C_x0000_00[2A:10]_dct[1:0][DqsRcvEnGrossDelay]$.
- 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.

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

The optimal values for $D18F2x214_dct[1:0][TwrwrSdSc, TwrwrSdDc, TwrwrDd]$ 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) = $MAX(TwrwrSdSc, CEIL(MAX(WOD + 3, CDD_{TwrwrSdDc} / 2 + (IF\ (D18F2x94_dct[1:0][FourRankRDimm1] | D18F2x94_dct[1:0][FourRankRDimm0])\ THEN\ 3.5\ ELSE\ 3\ ENDIF))))$.
- $TwrwrDd$ (in MEMCLKs) = $MAX(TwrwrSdDc, CEIL(MAX(WOD + 3, CDD_{TwrwrDd} / 2 + 3.5)))$.

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

- Each delay difference is $D18F2x9C_x0000_00[4A:30]_dct[1:0][WrDqsGrossDly]$ minus $D18F2x9C_x0000_00[4A:30]_dct[1:0][WrDqsGrossDly]$.
- 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.

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

The optimal value for $D18F2x218_dct[1:0][Twrrd]$ 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:

- IF (LRDIMM) THEN
 $Twrrd$ (in MEMCLKs) = $MAX(1, CEIL(MAX(WOD - BufDatDelay, CDD_{Twrrd} / 2 + 0.5 - WrEarly, (DdrRate \geq 1866 ? 1 : 0)) - LD + 3))$
 ELSE
 $Twrrd$ (in MEMCLKs) = $MAX(1, CEIL(MAX(WOD, CDD_{Twrrd} / 2 + 0.5 - WrEarly) - LD + 3))$
 ENDIF.

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

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

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

The optimal value for $D18F2x21C_dct[1:0][TrwtTO]$ 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) = $CEIL(MAX(ROD + BufDatDelay, CDD_{TrwtTO} / 2 - 0.5 + WrEarly) + LD + 3)$.
- If 1 DIMM/ch, substitute $ROD = 0$ and $BufDatDelay = 0$ in the above equation.

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

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

2.10.5.5.5 DRAM ODT Control

This section describes the ODT configurations and settings for the attached DIMMs. The tables specify ODT values for different speeds and configurations on a per channel basis. For UDIMMs and RDIMMs, the DIMM termination values are programmed to values as specified below via MRS commands. For LRDIMMs, the host-to-buffer termination values are programmed via RCW commands. See the SPD for the buffer-to-DRAM values and patterns to program via MRS commands. See [2.10.5.7 \[DRAM Device and Controller Initialization\]](#).

These tables document the optimal settings for motherboards which meet the relevant motherboard design guidelines. See [2.10 \[DRAM Controllers \(DCTs\)\]](#) for an overview of the DIMM population and memory bus speed support.

Table 56: BIOS Recommendations for RttNom and RttWr (G34r1 & UDIMM)

Condition								MR1 _dct[1:0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
1	667, 800	1.5, 1.35, 1.25	SR	-	-	SR	0	010b	00b
1	667, 800	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	010b	00b
1	1066, 1333	1.5, 1.35, 1.25	SR	-	-	SR	0	001b	00b
1	1066, 1333	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	001b	00b
1	1600	1.5, 1.35	SR	-	-	SR	0	011b	00b
1	1600	1.5, 1.35	DR	-	-	DR	0, 1	011b	00b
1	1866	1.5	SR	-	-	SR	0	011b	00b
1	1866	1.5	DR	-	-	DR	0, 1	011b	00b
2	667, 800	1.5, 1.35, 1.25	NP	SR	-	SR	0	010b	00b
2	667, 800	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	010b	00b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	SR, DR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	DR	-	DR	0, 1	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR, DR	-	DR	0, 1	011b	10b
2	1066, 1333	1.5, 1.35, 1.25	NP	SR	-	SR	0	001b	00b
2	1066, 1333	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	001b	00b
2	1333	1.5, 1.35	SR	SR, DR	-	SR	0	101b	10b
2	1333	1.5, 1.35	SR	DR	-	DR	0, 1	101b	10b
2	1333	1.5, 1.35	DR	SR	-	SR	0	101b	10b
2	1333	1.5, 1.35	DR	SR, DR	-	DR	0, 1	101b	10b
2	1333	1.25	SR	SR	-	SR	0	101b	10b
2	1600	1.5	NP	SR	-	SR	0	011b	00b
2	1600	1.5	NP	DR	-	DR	0, 1	011b	00b
2	1600	1.5	SR	SR	-	SR	0	100b	01b
3	667, 800, 1066	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	10b

Table 56: BIOS Recommendations for RttNom and RttWr (G34r1 & UDIMM)

Condition								MR1 _dct[1:0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
3	667, 800, 1066	1.5, 1.35, 1.25	SR	NP	SR	SR	0	011b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	SR	NP	DR	SR, DR	0	011b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	SR	NP	DR	DR	1	000b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	DR	NP	SR	SR, DR	0	011b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	DR	NP	SR, DR	DR	1	000b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	DR	NP	DR	DR	0	011b	10b
3	1333	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	01b
3	1333	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	01b
3	1333	1.5, 1.35, 1.25	SR	NP	SR	SR	0	101b	10b
3	1333	1.5, 1.35	SR	NP	DR	SR, DR	0	101b	10b
3	1333	1.5, 1.35	SR	NP	DR	DR	1	000b	10b
3	1333	1.5, 1.35	DR	NP	SR	SR, DR	0	101b	10b
3	1333	1.5, 1.35	DR	NP	SR, DR	DR	1	000b	10b
3	1333	1.5, 1.35	DR	NP	DR	DR	0	101b	10b
3	1600	1.5	NP	NP	SR	SR	0	000b	01b
3	1600	1.5	NP	NP	DR	DR	0, 1	000b	01b

Table 57: BIOS Recommendations for RttNom and RttWr (G34r1 & RDIMM)

Condition								MR1 _dct[1:0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
1	667, 800	1.5, 1.35, 1.25	SR	-	-	SR	0	010b	00b
1	667, 800	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	010b	00b
1	667, 800	1.5, 1.35, 1.25	QR	-	-	QR	0, 2	010b	10b
1	667, 800, 1066	1.5, 1.35, 1.25	QR	-	-	QR	1, 3	000b	10b
1	1066, 1333	1.5, 1.35, 1.25	SR	-	-	SR	0	001b	00b
1	1066, 1333	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	001b	00b
1	1066	1.5, 1.35, 1.25	QR	-	-	QR	0, 2	001b	10b
1	1333	1.5, 1.35	QR	-	-	QR	0, 2	011b	10b
1	1333	1.5, 1.35	QR	-	-	QR	1, 3	000b	10b
1	1600	1.5, 1.35	SR	-	-	SR	0	011b	00b
1	1600	1.5, 1.35	DR	-	-	DR	0, 1	011b	00b
1	1866	1.5	SR	-	-	SR	0	011b	00b
1	1866	1.5	DR	-	-	DR	0, 1	011b	00b
2	667, 800	1.5, 1.35, 1.25	NP	SR	-	SR	0	010b	00b
2	667, 800	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	010b	00b
2	667, 800	1.5, 1.35, 1.25	NP	QR	-	QR	0, 2	010b	10b
2	667, 800	1.5, 1.35, 1.25	NP, SR, DR	QR	-	QR	1, 3	000b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	SR, DR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	DR	-	DR	0, 1	011b	10b
2	667, 800	1.5, 1.35, 1.25	SR	QR	-	SR	0	101b	10b
2	667, 800	1.5, 1.35, 1.25	SR, DR	QR	-	QR	0, 2	001b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR, DR	-	DR	0, 1	011b	10b
2	667, 800	1.5, 1.35, 1.25	DR	QR	-	DR	0, 1	101b	10b
2	667, 800	1.5, 1.35, 1.25	QR	SR	-	SR	0	101b	10b
2	667, 800	1.5, 1.35, 1.25	QR	SR, DR, QR	-	QR	0, 2	001b	10b

Table 57: BIOS Recommendations for RttNom and RttWr (G34r1 & RDIMM)

Condition								MR1 _dct[1:0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
2	667, 800	1.5, 1.35, 1.25	QR	SR, DR, QR	-	QR	1, 3	000b	10b
2	667, 800	1.5, 1.35, 1.25	QR	DR	-	DR	0, 1	101b	10b
2	1066, 1333	1.5, 1.35, 1.25	NP	SR	-	SR	0	001b	00b
2	1066, 1333	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	001b	00b
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	0, 2	001b	10b
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	1, 3	000b	10b
2	1066	1.5	SR	QR	-	SR	0	101b	01b
2	1066	1.5	SR, DR	QR	-	QR	0, 2	001b	01b
2	1066	1.5	SR, DR	QR	-	QR	1, 3	000b	01b
2	1066	1.5	DR	QR	-	DR	0, 1	101b	01b
2	1066	1.5	QR	SR	-	SR	0	101b	01b
2	1066	1.5	QR	SR, DR, QR	-	QR	0, 2	001b	01b
2	1066	1.5	QR	SR, DR, QR	-	QR	1, 3	000b	01b
2	1066	1.5	QR	DR	-	DR	0, 1	101b	01b
2	1333	1.5	NP	QR	-	QR	0, 2	011b	10b
2	1333	1.5	NP	QR	-	QR	1, 3	000b	10b
2	1333	1.5, 1.35	SR	SR, DR	-	SR	0	101b	10b
2	1333	1.5, 1.35	SR	DR	-	DR	0, 1	101b	10b
2	1333	1.5, 1.35	DR	SR	-	SR	0	101b	10b
2	1333	1.5, 1.35	DR	SR, DR	-	DR	0, 1	101b	10b
2	1333	1.25	SR	SR	-	SR	0	101b	10b
2	1600	1.5	NP	SR	-	SR	0	011b	00b
2	1600	1.5	NP	DR	-	DR	0, 1	011b	00b
2	1600	1.5	SR	SR, DR	-	SR	0	100b	01b
2	1600	1.5	SR	DR	-	DR	0, 1	100b	01b
2	1600	1.5	DR	SR	-	SR	0	100b	01b
2	1600	1.5	DR	SR, DR	-	DR	0, 1	100b	01b
3	667, 800	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	10b
3	667, 800	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	10b
3	667	1.5, 1.35, 1.25	NP	QR	NP	QR	0, 2	010b	10b

Table 57: BIOS Recommendations for RttNom and RttWr (G34r1 & RDIMM)

Condition								MR1 _dct[1:0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
3	667, 800	1.5, 1.35, 1.25	NP	QR	NP, SR, DR	QR	1, 3	000b	10b
3	667	1.5, 1.35, 1.25	NP, SR	QR	SR	SR	0	101b	10b
3	667	1.5, 1.35, 1.25	NP, SR, DR	QR	SR, DR	QR	0, 2	001b	10b
3	667	1.5, 1.35, 1.25	NP, DR	QR	DR	DR	0	101b	10b
3	667, 800	1.5, 1.35, 1.25	NP	QR	DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	SR	NP, SR	SR	SR	0	011b	10b
3	667	1.5, 1.35, 1.25	SR	NP, SR	DR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	SR	NP, SR, QR	DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	SR	DR	SR, DR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	SR	DR	SR, DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	SR, DR	QR	SR, DR	QR	1, 3	000b	10b
3	667	1.5, 1.35, 1.25	SR	QR	DR	SR, DR	0	101b	10b
3	667	1.5, 1.35, 1.25	DR	NP, DR	SR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	DR	NP, SR, DR, QR	SR, DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	DR	NP, DR	DR	DR	0	011b	10b
3	667	1.5, 1.35, 1.25	DR	SR	SR, DR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	DR	QR	SR	SR, DR	0	101b	10b
3	800	1.5, 1.35, 1.25	NP	QR	NP, SR, DR	QR	0, 2	001b	10b
3	800	1.5, 1.35, 1.25	NP	QR	SR	SR	0	101b	10b
3	800	1.5, 1.35, 1.25	NP	QR	DR	DR	0	101b	10b
3	800, 1066	1.5, 1.35, 1.25	SR	NP	SR	SR	0	011b	10b
3	800, 1066	1.5, 1.35, 1.25	SR	NP	DR	SR, DR	0	011b	10b
3	800	1.5, 1.35	SR	NP, SR, QR	DR	DR	1	000b	10b
3	800	1.5, 1.35	SR	SR, QR	SR	SR	0	101b	10b
3	800	1.5, 1.35	SR	SR, QR	DR	SR, DR	0	101b	10b
3	800	1.5, 1.35	SR	DR	SR, DR	SR, DR	0	101b	10b
3	800	1.5, 1.35	SR	DR	SR, DR	DR	1	000b	10b
3	800	1.5, 1.35	SR, DR	QR	SR, DR	QR	0, 2	001b	10b

Table 57: BIOS Recommendations for RttNom and RttWr (G34r1 & RDIMM)

Condition								MR1 _dct[1:0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
3	800	1.5, 1.35	SR, DR	QR	SR, DR	QR	1, 3	000b	10b
3	800, 1066	1.5, 1.35, 1.25	DR	NP	SR	SR, DR	0	011b	10b
3	800	1.5, 1.35	DR	NP, SR, DR, QR	SR, DR	DR	1	000b	10b
3	800, 1066	1.5, 1.35, 1.25	DR	NP	DR	DR	0	011b	10b
3	800	1.5, 1.35	DR	SR	SR, DR	SR, DR	0	101b	10b
3	800	1.5, 1.35	DR	DR, QR	SR	SR, DR	0	101b	10b
3	800	1.5, 1.35	DR	DR, QR	DR	DR	0	101b	10b
3	800	1.25	SR	SR	SR	SR	0	101b	10b
3	800	1.25	SR	NP	DR	DR	1	000b	10b
3	800	1.25	DR	NP	SR, DR	DR	1	000b	10b
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	01b
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	01b
3	1066	1.5, 1.35	NP	QR	NP	QR	0, 2	001b	10b
3	1066	1.5, 1.35	NP	QR	NP	QR	1, 3	000b	10b
3	1066	1.5	SR	NP, SR	DR	DR	1	000b	10b
3	1066	1.5, 1.35	SR	SR	SR	SR	0	101b	10b
3	1066	1.5	SR	SR	DR	SR, DR	0	101b	10b
3	1066	1.5	SR	DR	SR, DR	SR, DR	0	101b	10b
3	1066	1.5	SR	DR	SR, DR	DR	1	000b	10b
3	1066	1.5	DR	NP, SR, DR	SR, DR	DR	1	000b	10b
3	1066	1.5	DR	SR	SR, DR	SR, DR	0	101b	10b
3	1066	1.5	DR	DR	SR	SR, DR	0	101b	10b
3	1066	1.5	DR	DR	DR	DR	0	101b	10b
3	1066	1.35, 1.25	SR	NP	DR	DR	1	000b	10b
3	1066	1.35, 1.25	DR	NP	SR, DR	DR	1	000b	10b
3	1333	1.5, 1.35, 1.25	SR	NP	SR	SR	0	101b	10b
3	1333	1.5, 1.35	SR	NP	DR	SR, DR	0	101b	10b
3	1333	1.5, 1.35	SR	NP	DR	DR	1	000b	10b
3	1333	1.5, 1.35	DR	NP	SR	SR, DR	0	101b	10b
3	1333	1.5, 1.35	DR	NP	SR, DR	DR	1	000b	10b

Table 57: BIOS Recommendations for RttNom and RttWr (G34r1 & RDIMM)

Condition								MR1 _dct[1:0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
3	1333	1.5, 1.35	DR	NP	DR	DR	0	101b	10b
3	1600	1.5	NP	NP	SR	SR	0	000b	01b
3	1600	1.5	NP	NP	DR	DR	0, 1	000b	01b

Table 58: BIOS Recommendations for RttNom and RttWr (G34r1 & LRDIMM)

Condition						F3RC0	F3RC1
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	RttNom	RttWr
1	667, 800	1.5, 1.35, 1.25	LR	-	-	010b	000b
1	1066, 1333	1.5, 1.35, 1.25	LR	-	-	001b	000b
1	1600	1.5, 1.35	LR	-	-	011b	000b
1	1866	1.5	LR	-	-	011b	000b
2	667, 800	1.5, 1.35, 1.25	NP	LR	-	010b	000b
2	667, 800, 1066	1.5, 1.35, 1.25	LR	LR	-	011b	010b
2	1066, 1333	1.5, 1.35, 1.25	NP	LR	-	001b	000b
2	1333	1.5, 1.35	LR	LR	-	101b	010b
2	1600	1.5	NP	LR	-	011b	000b
2	1600	1.5	LR	LR	-	101b	001b
3	667, 800	1.5, 1.35, 1.25	NP	NP	LR	000b	010b
3	667, 800	1.5, 1.35, 1.25	LR	NP, LR	LR	011b	010b
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	LR	000b	001b
3	1066	1.5, 1.35	LR	NP, LR	LR	011b	010b
3	1066	1.25	LR	NP	LR	011b	010b
3	1333	1.5	LR	NP, LR	LR	101b	010b
3	1333	1.35	LR	NP	LR	101b	010b
3	1600	1.5	NP	NP	LR	000b	001b
3	1600	1.5	LR	NP	LR	101b	001b

Table 59: BIOS Recommendations for RttNom and RttWr (C32r1 & UDIMM)

Condition								MR1_ _dct[1: 0]	MR2_ _dct[1: 1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
1	667, 800	1.5, 1.35, 1.25	SR	-	-	SR	0	010b	00b
1	667, 800	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	010b	00b
1	1066, 1333	1.5, 1.35, 1.25	SR	-	-	SR	0	001b	00b
1	1066, 1333	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	001b	00b
1	1600	1.5	SR	-	-	SR	0	011b	00b
1	1600	1.5	DR	-	-	DR	0, 1	011b	00b
2	667, 800	1.5, 1.35, 1.25	NP	SR	-	SR	0	010b	00b
2	667, 800	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	010b	00b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	SR, DR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	DR	-	DR	0, 1	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR, DR	-	DR	0, 1	011b	10b
2	1066, 1333	1.5, 1.35, 1.25	NP	SR	-	SR	0	001b	00b
2	1066, 1333	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	001b	00b
2	1333	1.5, 1.35	SR	SR	-	SR	0	101b	10b
2	1600	1.5	NP	SR	-	SR	0	011b	00b
2	1600	1.5	NP	DR	-	DR	0, 1	011b	00b
3	667, 800, 1066	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	SR	NP	SR	SR	0	011b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	SR	NP	DR	SR, DR	0	011b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	SR	NP	DR	DR	1	000b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	DR	NP	SR	SR, DR	0	011b	10b

Table 59: BIOS Recommendations for RttNom and RttWr (C32r1 & UDIMM)

Condition								MR1_ dct[1: 0]	MR2_ _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
3	667, 800, 1066	1.5, 1.35, 1.25	DR	NP	SR, DR	DR	1	000b	10b
3	667, 800, 1066	1.5, 1.35, 1.25	DR	NP	DR	DR	0	011b	10b
3	1333	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	01b
3	1333	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	01b
3	1333	1.5, 1.35	SR	NP	SR	SR	0	101b	10b
3	1600	1.5	NP	NP	SR	SR	0	000b	01b
3	1600	1.5	NP	NP	DR	DR	0, 1	000b	01b

Table 60: BIOS Recommendations for RttNom and RttWr (C32r1 & RDIMM)

Condition								MR1_ _dct[1: 0]	MR2_ _dct[1: 1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
1	667, 800	1.5, 1.35, 1.25	SR	-	-	SR	0	010b	00b
1	667, 800	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	010b	00b
1	667, 800	1.5, 1.35, 1.25	QR	-	-	QR	0, 2	010b	10b
1	667, 800, 1066	1.5, 1.35, 1.25	QR	-	-	QR	1, 3	000b	10b
1	1066, 1333	1.5, 1.35, 1.25	SR	-	-	SR	0	001b	00b
1	1066, 1333	1.5, 1.35, 1.25	DR	-	-	DR	0, 1	001b	00b
1	1066	1.5, 1.35, 1.25	QR	-	-	QR	0, 2	001b	10b
1	1333	1.5	QR	-	-	QR	0, 2	011b	10b
1	1333	1.5	QR	-	-	QR	1, 3	000b	10b
1	1600	1.5	SR	-	-	SR	0	011b	00b
1	1600	1.5	DR	-	-	DR	0, 1	011b	00b
2	667, 800	1.5, 1.35, 1.25	NP	SR	-	SR	0	010b	00b
2	667, 800	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	010b	00b
2	667, 800	1.5, 1.35, 1.25	NP	QR	-	QR	0, 2	010b	10b
2	667, 800	1.5, 1.35, 1.25	NP, SR, DR	QR	-	QR	1, 3	000b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	SR, DR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	SR	DR	-	DR	0, 1	011b	10b
2	667, 800	1.5, 1.35, 1.25	SR	QR	-	SR	0	101b	10b
2	667, 800	1.5, 1.35, 1.25	SR, DR	QR	-	QR	0, 2	001b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR	-	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR, DR	-	DR	0, 1	011b	10b
2	667, 800	1.5, 1.35, 1.25	DR	QR	-	DR	0, 1	101b	10b
2	667, 800	1.5, 1.35, 1.25	QR	SR	-	SR	0	101b	10b
2	667, 800	1.5, 1.35, 1.25	QR	SR, DR, QR	-	QR	0, 2	001b	10b
2	667, 800	1.5, 1.35, 1.25	QR	SR, DR, QR	-	QR	1, 3	000b	10b
2	667, 800	1.5, 1.35, 1.25	QR	DR	-	DR	0, 1	101b	10b

Table 60: BIOS Recommendations for RttNom and RttWr (C32r1 & RDIMM)

Condition								MR1_ dct[1: 0]	MR2_ _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
2	1066, 1333	1.5, 1.35, 1.25	NP	SR	-	SR	0	001b	00b
2	1066, 1333	1.5, 1.35, 1.25	NP	DR	-	DR	0, 1	001b	00b
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	0, 2	001b	10b
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	1, 3	000b	10b
2	1333	1.5	NP	QR	-	QR	0, 2	011b	10b
2	1333	1.5	NP	QR	-	QR	1, 3	000b	10b
2	1333	1.5, 1.35	SR	SR	-	SR	0	101b	10b
2	1600	1.5	NP	SR	-	SR	0	011b	00b
2	1600	1.5	NP	DR	-	DR	0, 1	011b	00b
3	667, 800	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	10b
3	667, 800	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	10b
3	667	1.5, 1.35, 1.25	NP	QR	NP	QR	0, 2	010b	10b
3	667, 800	1.5, 1.35, 1.25	NP	QR	NP, SR, DR	QR	1, 3	000b	10b
3	667	1.5, 1.35, 1.25	NP, SR	QR	SR	SR	0	101b	10b
3	667	1.5, 1.35, 1.25	NP, SR, DR	QR	SR, DR	QR	0, 2	001b	10b
3	667	1.5, 1.35, 1.25	NP, DR	QR	DR	DR	0	101b	10b
3	667, 800	1.5, 1.35, 1.25	NP	QR	DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	SR	NP, SR	SR	SR	0	011b	10b
3	667	1.5, 1.35, 1.25	SR	NP, SR	DR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	SR	NP, SR, QR	DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	SR	DR	SR, DR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	SR	DR	SR, DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	SR, DR	QR	SR, DR	QR	1, 3	000b	10b
3	667	1.5, 1.35, 1.25	SR	QR	DR	SR, DR	0	101b	10b
3	667	1.5, 1.35, 1.25	DR	NP, DR	SR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	DR	NP, SR, DR, QR	SR, DR	DR	1	000b	10b
3	667	1.5, 1.35, 1.25	DR	NP, DR	DR	DR	0	011b	10b
3	667	1.5, 1.35, 1.25	DR	SR	SR, DR	SR, DR	0	011b	10b
3	667	1.5, 1.35, 1.25	DR	QR	SR	SR, DR	0	101b	10b

Table 60: BIOS Recommendations for RttNom and RttWr (C32r1 & RDIMM)

Condition								MR1_ dct[1: 0]	MR2_ dct[1: 1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	Rank	RttNom[2:0]	RttWr
3	800	1.5, 1.35, 1.25	NP	QR	NP, SR, DR	QR	0, 2	001b	10b
3	800	1.5, 1.35, 1.25	NP	QR	SR	SR	0	101b	10b
3	800	1.5, 1.35, 1.25	NP	QR	DR	DR	0	101b	10b
3	800, 1066	1.5, 1.35, 1.25	SR	NP	SR	SR	0	011b	10b
3	800, 1066	1.5, 1.35, 1.25	SR	NP	DR	SR, DR	0	011b	10b
3	800	1.5, 1.35	SR	NP, SR	DR	DR	1	000b	10b
3	800	1.5, 1.35, 1.25	SR	SR	SR	SR	0	101b	10b
3	800	1.5, 1.35	SR	SR	DR	SR, DR	0	101b	10b
3	800	1.5, 1.35	SR	DR	SR, DR	SR, DR	0	101b	10b
3	800	1.5, 1.35	SR	DR	SR, DR	DR	1	000b	10b
3	800, 1066	1.5, 1.35, 1.25	DR	NP	SR	SR, DR	0	011b	10b
3	800	1.5, 1.35	DR	NP, SR, DR	SR, DR	DR	1	000b	10b
3	800, 1066	1.5, 1.35, 1.25	DR	NP	DR	DR	0	011b	10b
3	800	1.5, 1.35	DR	SR	SR, DR	SR, DR	0	101b	10b
3	800	1.5, 1.35	DR	DR	SR	SR, DR	0	101b	10b
3	800	1.5, 1.35	DR	DR	DR	DR	0	101b	10b
3	800	1.25	SR	NP	DR	DR	1	000b	10b
3	800	1.25	DR	NP	SR, DR	DR	1	000b	10b
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	SR	SR	0	000b	01b
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	DR	DR	0, 1	000b	01b
3	1066	1.5	NP	QR	NP	QR	0, 2	001b	10b
3	1066	1.5	NP	QR	NP	QR	1, 3	000b	10b
3	1066	1.5, 1.35, 1.25	SR	NP	DR	DR	1	000b	10b
3	1066	1.5	SR	SR	SR	SR	0	101b	10b
3	1066	1.5, 1.35, 1.25	DR	NP	SR, DR	DR	1	000b	10b
3	1333	1.5, 1.35	SR	NP	SR	SR	0	101b	10b
3	1600	1.5	NP	NP	SR	SR	0	000b	01b
3	1600	1.5	NP	NP	DR	DR	0, 1	000b	01b

Table 61: BIOS Recommendations for RttNom and RttWr (C32r1 & LRDIMM)

Condition						F3RC0	F3RC1
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	RttNom	RttWr
						1	667, 800
1	1066, 1333	1.5, 1.35, 1.25	LR	-	-	001b	000b
1	1600	1.5, 1.35	LR	-	-	011b	000b
1	1866	1.5	LR	-	-	011b	000b
2	667, 800	1.5, 1.35, 1.25	NP	LR	-	010b	000b
2	667, 800, 1066	1.5, 1.35, 1.25	LR	LR	-	011b	010b
2	1066, 1333	1.5, 1.35, 1.25	NP	LR	-	001b	000b
2	1333	1.5, 1.35	LR	LR	-	101b	010b
2	1600	1.5	NP	LR	-	011b	000b
3	667, 800	1.5, 1.35, 1.25	NP	NP	LR	000b	010b
3	667, 800	1.5, 1.35, 1.25	LR	NP, LR	LR	011b	010b
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	LR	000b	001b
3	1066	1.5	LR	NP, LR	LR	011b	010b
3	1066	1.35, 1.25	LR	NP	LR	011b	010b
3	1333	1.5, 1.35	LR	NP	LR	101b	010b
3	1600	1.5	NP	NP	LR	000b	001b

Table 62: BIOS Recommendations for RttNom and RttWr (AM3r2 & UDIMM)

Condition							MR1_ dct[1: 0]	MR2_ _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM	Rank	RttNom[2:0]	RttWr
							1	667, 800, 1066
1	667, 800, 1066	1.5, 1.35	DR	-	DR	0, 1	010b	00b
1	1333	1.5, 1.35	SR	-	SR	0	001b	00b
1	1333	1.5, 1.35	DR	-	DR	0, 1	001b	00b
1	1600, 1866	1.5	SR	-	SR	0	011b	00b
1	1600, 1866	1.5	DR	-	DR	0, 1	011b	00b
2	667, 800, 1066	1.5, 1.35	NP	SR	SR	0	010b	00b
2	667, 800, 1066	1.5, 1.35	NP	DR	DR	0, 1	010b	00b

Table 62: BIOS Recommendations for RttNom and RttWr (AM3r2 & UDIMM)

Condition							MR1_ dct[1: 0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM	Rank	RttNom[2:0]	RttWr
2	667, 800, 1066	1.5, 1.35	SR	SR, DR	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35	SR	DR	DR	0, 1	011b	10b
2	667, 800, 1066	1.5, 1.35	DR	SR	SR	0	011b	10b
2	667, 800, 1066	1.5, 1.35	DR	SR, DR	DR	0, 1	011b	10b
2	1333	1.5, 1.35	NP	SR	SR	0	001b	00b
2	1333	1.5, 1.35	NP	DR	DR	0, 1	001b	00b
2	1333	1.5	SR	SR, DR	SR	0	101b	10b
2	1333	1.5	SR	DR	DR	0, 1	101b	10b
2	1333	1.5	DR	SR	SR	0	101b	10b
2	1333	1.5	DR	SR, DR	DR	0, 1	101b	10b
2	1333	1.35	SR	SR	SR	0	101b	10b
2	1600	1.5	NP	SR	SR	0	011b	00b
2	1600	1.5	NP	DR	DR	0, 1	011b	00b
2	1600	1.5	SR	SR	SR	0	100b	01b

Table 63: BIOS Recommendations for RttNom and RttWr (SODIMM)

Condition							MR1_ dct[1: 0]	MR2 _dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM	Rank	RttNom[2:0]	RttWr
1	667, 800, 1066	1.5, 1.35	SR	-	SR	0	010b	00b
1	667, 800, 1066	1.5, 1.35	DR	-	DR	0, 1	010b	00b
1	1333	1.5, 1.35	SR	-	SR	0	001b	00b
1	1333	1.5, 1.35	DR	-	DR	0, 1	001b	00b
1	1600	1.5	SR	-	SR	0	001b	00b
1	1600	1.5	DR	-	DR	0, 1	001b	00b
1	1600	1.35	SR	-	SR	0	001b	00b

Table 63: BIOS Recommendations for RttNom and RttWr (SODIMM)

Condition							MR1_ dct[1: 0]	MR2_ dct[1: 1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM	Rank	RttNom[2:0]	RttWr
1	1600	1.35	DR	-	DR	0, 1	001b	00b
2	667, 800, 1066	1.5, 1.35	NP	SR	SR	0	010b	00b
2	667, 800, 1066	1.5, 1.35	NP	DR	DR	0, 1	010b	00b
2	667, 800	1.5, 1.35	SR	SR, DR	SR	0	011b	10b
2	667, 800	1.5, 1.35	SR	DR	DR	0, 1	011b	10b
2	667, 800	1.5, 1.35	DR	SR	SR	0	011b	10b
2	667, 800	1.5, 1.35	DR	SR, DR	DR	0, 1	011b	10b
2	1066	1.5, 1.35	SR	SR, DR	SR	0	101b	10b
2	1066	1.5, 1.35	SR	DR	DR	0, 1	101b	10b
2	1066	1.5, 1.35	DR	SR	SR	0	101b	10b
2	1066	1.5, 1.35	DR	SR, DR	DR	0, 1	101b	10b
2	1333	1.5, 1.35	NP	SR	SR	0	001b	00b
2	1333	1.5, 1.35	NP	DR	DR	0, 1	001b	00b
2	1333	1.5	SR	SR, DR	SR	0	101b	10b
2	1333	1.5	SR	DR	DR	0, 1	101b	10b
2	1333	1.5	DR	SR	SR	0	101b	10b
2	1333	1.5	DR	SR, DR	DR	0, 1	101b	10b
2	1333	1.35	SR	SR	SR	0	101b	10b

BIOS configures the DIMM ODT behavior per chip select according to the DIMM population. The ODT patterns for reads and writes are programmed using [D18F2x\[234:230\]_dct\[1:0\]](#) and [D18F2x\[23C:238\]_dct\[1:0\]](#), respectively, per the recommended values specified by the following tables. In all cases, the processor ODT is off for writes and is on for reads.

BIOS also configures the DIMM ODT pattern used during write leveling by setting [D18F2x9C_x0000_0008_dct\[1:0\]\[WrLvOdtEn\]](#) and programming [D18F2x9C_x0000_0008_dct\[1:0\]\[WrLvOdt\]](#). BIOS programs [D18F2x9C_x0000_0008_dct\[1:0\]\[WrLvOdt\]](#) with the [D18F2x\[23C:238\]_dct\[1:0\]](#) value provided for writes to the rank targeted by training. See 2.10.5.8.1 [Write Levelization Training].

See 2.10.2 [DDR Pad to Processor Pin Mapping] when programming [D18F2x\[234:230\]_dct\[1:0\]](#), [D18F2x\[23C:238\]_dct\[1:0\]](#), and [D18F2x9C_x0000_0008_dct\[1:0\]\[WrLvOdt\]](#).

Table 64: DDR3 UDIMM/SODIMM ODT Pattern 1 DIMM Slot/Ch

DIMM0 ¹	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
	D18F2x234	D18F2x230	D18F2x23C	D18F2x238
SR	0000_0000h	0000_0000h	0000_0000h	0000_0001h
DR	0000_0000h	0000_0000h	0000_0000h	0000_0401h

Table 65: DDR3 UDIMM/SODIMM ODT Pattern 2 DIMM Slots/Ch

DIMM0	DIMM1	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
		D18F2x234	D18F2x230	D18F2x23C	D18F2x238
-	SR	0000_0000h	0000_0000h	0000_0000h	0002_0000h
-	DR	0000_0000h	0000_0000h	0000_0000h	0802_0000h
SR/DR	SR/DR	0000_0000h	0101_0202h	0000_0000h	0903_0603h

Table 66: DDR3 UDIMM ODT Pattern 3 DIMM Slots/Ch

DIMM0	DIMM1	DIMM2	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
			D18F2x234	D18F2x230	D18F2x23C	D18F2x238
-	-	SR/DR	0000_0000h	0000_0000h	0000_0004h	0000_0000h
SR/DR	-	SR/DR	0000_0101h	0000_0404h	0000_0105h	0000_0405h

Table 67: DDR3 RDIMM ODT Pattern 1 DIMM Slot/Ch

DIMM0	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
	D18F2x234	D18F2x230	D18F2x23C	D18F2x238
SR	0000_0000h	0000_0000h	0000_0000h	0000_0001h
DR	0000_0000h	0000_0000h	0000_0000h	0000_0401h
QR	0000_0000h	0000_0000h	0000_0505h	0000_0505h

Table 68: DDR3 RDIMM ODT Pattern 2 DIMM Slots/Ch

DIMM0	DIMM1	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
		D18F2x234	D18F2x230	D18F2x23C	D18F2x238
-	SR	0000_0000h	0000_0000h	0000_0000h	0002_0000h
-	DR	0000_0000h	0000_0000h	0000_0000h	0802_0000h
-	QR	0000_0000h	0000_0000h	020A_0000h	080A_0000h
SR/DR	SR/DR	0000_0000h	0101_0202h	0000_0000h	0903_0603h
SR/DR	QR	0101_0000h	0101_0A0Ah	0109_0000h	0103_0E0Bh
QR	SR/DR	0000_0202h	0505_0202h	0000_0206h	0D07_0203h
QR	QR	0505_0A0Ah	0505_0A0Ah	050D_0A0Eh	0507_0A0Bh

Table 69: DDR3 RDIMM ODT Pattern 3 DIMM Slots/Ch

DIMM0	DIMM1	DIMM2	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
			D18F2x234	D18F2x230	D18F2x23C	D18F2x238
-	-	SR/DR	0000_0000h	0000_0000h	0000_0004h	0000_0000h
SR/DR	-	SR/DR	0000_0101h	0000_0404h	0000_0105h	0000_0405h
SR/DR	SR/DR	SR/DR	0000_0303h	0505_0606h	0000_0307h	0D07_0607h
-	QR	-	0000_0000h	0000_0000h	020A_0000h	080A_0000h
-	QR	SR/DR	0404_0A0Ah	0404_0000h	040C_0A0Eh	0406_0000h
SR/DR	QR	SR/DR	0505_0B0Bh	0505_0E0Eh	050D_0B0Fh	0507_0E0Fh

Table 70: DDR3 LRDIMM ODT Pattern 1 DIMM Slot/Ch

DIMM0	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
	D18F2x234	D18F2x230	D18F2x23C	D18F2x238
LR	0000_0000h	0000_0000h	0000_0101h	0000_0101h

Table 71: DDR3 LRDIMM ODT Pattern 2 DIMM Slots/Ch

DIMM0	DIMM1	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
		D18F2x234	D18F2x230	D18F2x23C	D18F2x238
-	LR	0000_0000h	0000_0000h	0202_0000h	0202_0000h
LR	LR	0101_0202h	0101_0202h	0303_0303h	0303_0303h

Table 72: DDR3 LRDIMM ODT Pattern 3 DIMM Slots/Ch

DIMM0	DIMM1	DIMM2	D18F2x[234:230]_dct[1:0]		D18F2x[23C:238]_dct[1:0]	
			D18F2x234	D18F2x230	D18F2x23C	D18F2x238
-	-	LR	0000_0000h	0000_0000h	0000_0404h	0000_0000h
LR	-	LR	0000_0101h	0000_0404h	0000_0505h	0000_0505h
LR	LR	LR	0000_0303h	0505_0606h	0000_0707h	0707_0707h

2.10.5.5.6 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. These tables document the optimal settings for motherboards which meet the relevant motherboard design guidelines. See [2.10 \[DRAM Controllers \(DCTs\)\]](#) for an overview of the DIMM population and memory bus speed support.

Table 73: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & UDIMM)

Condition						D18F2x94_dct[1:0]	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
1	667	1.5, 1.35, 1.25	SR	-	-	0	00000000h	00112222h
1	667	1.5, 1.35, 1.25	DR	-	-	0	003B0000h	00112222h
1	800	1.5, 1.35, 1.25	SR	-	-	0	00000000h	10112222h
1	800	1.5, 1.35, 1.25	DR	-	-	0	003B0000h	10112222h
1	1066	1.5, 1.35, 1.25	SR, DR	-	-	0	00383837h	20112222h
1	1333	1.5, 1.35, 1.25	SR, DR	-	-	0	00363635h	30112222h
1	1600	1.5, 1.35	SR	-	-	0	00353533h	30112222h
1	1600	1.5, 1.35	DR	-	-	1	00003533h	30112222h
1	1866	1.5	SR	-	-	0	00333330h	30332222h
1	1866	1.5	DR	-	-	1	00003330h	30332222h
2	667	1.5, 1.35, 1.25	NP	SR	-	0	00000000h	00112222h
2	667	1.5, 1.35, 1.25	NP	DR	-	0	003B0000h	00112222h
2	667	1.5, 1.35, 1.25	SR, DR	SR, DR	-	0	00390039h	10222222h
2	800	1.5, 1.35, 1.25	NP	SR	-	0	00000000h	10112222h
2	800	1.5, 1.35, 1.25	NP	DR	-	0	003B0000h	10112222h
2	800	1.5, 1.35, 1.25	SR, DR	SR, DR	-	0	00390039h	20222222h
2	1066	1.5, 1.35, 1.25	NP	SR, DR	-	0	00383837h	20112222h
2	1066	1.5, 1.35, 1.25	SR, DR	SR, DR	-	0	003A3A3Ah	30222222h
2	1333	1.5, 1.35, 1.25	NP	SR, DR	-	0	00363635h	30112222h
2	1333	1.5, 1.35, 1.25	SR	SR	-	1	00003939h	30222222h
2	1333	1.5, 1.35	SR	DR	-	1	00003938h	30222222h
2	1333	1.5, 1.35	DR	SR, DR	-	1	00003938h	30222222h
2	1600	1.5	NP	SR	-	0	00353533h	30112222h
2	1600	1.5	NP	DR	-	1	00003533h	30112222h
2	1600	1.5	SR	SR	-	1	00003738h	30222222h
3	667	1.5, 1.35, 1.25	NP	NP	SR	0	00000000h	00332222h
3	667	1.5, 1.35, 1.25	NP	NP	DR	0	003B0000h	00332222h
3	667	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00390039h	10222222h
3	800	1.5, 1.35, 1.25	NP	NP	SR	0	00000000h	10332222h

Table 73: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & UDIMM)

Condition						D18F2x94_dctl[1:0]	SlowAccessMode	D18F2x9C_x0000_0004_dctl[1:0]	D18F2x9C_x0000_0000_dctl[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2				
3	800	1.5, 1.35, 1.25	NP	NP	DR	0	003B0000h	10332222h	
3	800	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00390039h	20222222h	
3	1066	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00383837h	20332222h	
3	1066	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	003A3A3Ah	30222222h	
3	1333	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00363635h	30332222h	
3	1333	1.5, 1.35, 1.25	SR	NP	SR	1	00003939h	30222222h	
3	1333	1.5, 1.35	SR	NP	DR	1	00003938h	30222222h	
3	1333	1.5, 1.35	DR	NP	SR, DR	1	00003938h	30222222h	
3	1600	1.5	NP	NP	SR	0	00353533h	30332222h	
3	1600	1.5	NP	NP	DR	1	00003533h	30332222h	

Table 74: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & RDIMM)

Condition						D18F2x94_dctl[1:0]	SlowAccessMode	D18F2x9C_x0000_0004_dctl[1:0]	D18F2x9C_x0000_0000_dctl[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2				
1	667	1.5, 1.35, 1.25	SR, DR	-	-	0	00000000h	00112222h	
1	667	1.5, 1.35, 1.25	QR	-	-	0	00000000h	00222222h	
1	800	1.5, 1.35, 1.25	SR, DR	-	-	0	00000000h	10112222h	
1	800	1.5, 1.35, 1.25	QR	-	-	0	00000000h	10222222h	
1	1066	1.5, 1.35, 1.25	SR, DR	-	-	0	003C3C3Ch	20112222h	
1	1066	1.5, 1.35, 1.25	QR	-	-	0	003C3C3Ch	30222222h	

Table 74: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & RDIMM)

Condition						D18F2x94_dctl[1:0] SlowAccessMode	D18F2x9C_x0000_0004_dctl[1:0]	D18F2x9C_x0000_0000_dctl[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
1	1333	1.5, 1.35, 1.25	SR, DR	-	-	0	003A3A3Ah	30112222h
1	1333	1.5, 1.35	QR	-	-	0	003A3A3Ah	30222222h
1	1600	1.5, 1.35	SR, DR	-	-	0	00393939h	30112222h
1	1866	1.5	SR, DR	-	-	0	00393939h	30332222h
2	667	1.5, 1.35, 1.25	NP	SR, DR	-	0	00000000h	00112222h
2	667	1.5, 1.35, 1.25	NP	QR	-	0	00000000h	00222222h
2	667	1.5, 1.35, 1.25	SR, DR, QR	SR, DR, QR	-	0	00000000h	10222222h
2	800	1.5, 1.35, 1.25	NP	SR, DR	-	0	00000000h	10112222h
2	800	1.5, 1.35, 1.25	NP	QR	-	0	00000000h	10222222h
2	800	1.5, 1.35, 1.25	SR, DR, QR	SR, DR, QR	-	0	00000000h	20222222h
2	1066	1.5, 1.35, 1.25	NP	SR, DR	-	0	00393C39h	20112222h
2	1066	1.5, 1.35, 1.25	NP	QR	-	0	00393C39h	20222222h
2	1066	1.5	SR, DR, QR	SR, DR, QR	-	0	003A3C3Ah	30222222h
2	1066	1.35, 1.25	SR, DR	SR, DR	-	0	003A3C3Ah	30222222h
2	1333	1.5, 1.35, 1.25	NP	SR, DR	-	0	00373A37h	30112222h
2	1333	1.5	NP	QR	-	0	00373A37h	30222222h
2	1333	1.5, 1.35	SR, DR	SR, DR	-	0	00383A38h	30222222h
2	1333	1.25	SR	SR	-	0	00383A38h	30222222h
2	1600	1.5	NP	SR, DR	-	0	00363936h	30112222h
2	1600	1.5	SR, DR	SR, DR	-	0	00353935h	30222222h
3	667	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00000000h	00332222h
3	667	1.5, 1.35, 1.25	NP	QR	NP	0	00000000h	10222222h
3	667	1.5, 1.35, 1.25	NP	QR	SR, DR	0	00000000h	20222222h
3	667	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00000000h	10222222h
3	667	1.5, 1.35, 1.25	SR, DR	SR, DR, QR	SR, DR	0	00380038h	30112222h
3	800	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00390039h	10332222h

Table 74: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & RDIMM)

Condition						D18F2x94_dctl[1:0]	D18F2x9C_x0000_0004_dctl[1:0]	D18F2x9C_x0000_0000_dctl[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	SlowAccessMode		
3	800	1.5, 1.35, 1.25	NP	QR	NP	0	00390039h	20222222h
3	800	1.5, 1.35, 1.25	NP	QR	SR, DR	0	003A003Ah	30222222h
3	800	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	003A003Ah	20222222h
3	800	1.5, 1.35	SR, DR	SR, DR, QR	SR, DR	0	00360036h	30112222h
3	800	1.25	SR	SR	SR	0	00360036h	30112222h
3	1066	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00373C37h	20332222h
3	1066	1.5, 1.35	NP	QR	NP	0	00373C37h	30222222h
3	1066	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00383C38h	30222222h
3	1066	1.5	SR, DR	SR, DR	SR, DR	0	00333C33h	30112222h
3	1066	1.35	SR	SR	SR	0	00333C33h	30112222h
3	1333	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00353A35h	30332222h
3	1333	1.5, 1.35	SR, DR	NP	SR, DR	0	00363A36h	30222222h
3	1333	1.25	SR	NP	SR	0	00363A36h	30222222h
3	1600	1.5	NP	NP	SR, DR	0	00333933h	30332222h

Table 75: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & LRDIMM)

Condition						D18F2x94_dct[1:0]	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	SlowAccessMode		
1	667	1.5, 1.35, 1.25	LR	-	-	0	00000000h	00112222h
1	800	1.5, 1.35, 1.25	LR	-	-	0	00000000h	10112222h
1	1066	1.5, 1.35, 1.25	LR	-	-	0	003C3C3Ch	20112222h
1	1333	1.5, 1.35, 1.25	LR	-	-	0	003A3A3Ah	30112222h
1	1600	1.5, 1.35	LR	-	-	0	00393939h	30112222h
1	1866	1.5	LR	-	-	0	00393939h	30332222h
2	667	1.5, 1.35, 1.25	NP	LR	-	0	00000000h	00112222h
2	667	1.5, 1.35, 1.25	LR	LR	-	0	00000000h	10222222h
2	800	1.5, 1.35, 1.25	NP	LR	-	0	00000000h	10112222h
2	800	1.5, 1.35, 1.25	LR	LR	-	0	00000000h	20222222h
2	1066	1.5, 1.35, 1.25	NP	LR	-	0	00393C39h	20112222h
2	1066	1.5, 1.35, 1.25	LR	LR	-	0	003A3C3Ah	30222222h
2	1333	1.5, 1.35, 1.25	NP	LR	-	0	00373A37h	30112222h
2	1333	1.5, 1.35	LR	LR	-	0	00383A38h	30222222h
2	1600	1.5	NP	LR	-	0	00363936h	30112222h
2	1600	1.5	LR	LR	-	0	00353935h	30222222h
3	667	1.5, 1.35, 1.25	NP	NP	LR	0	00000000h	00332222h
3	667	1.5, 1.35, 1.25	LR	NP	LR	0	00000000h	20222222h
3	667	1.5, 1.35, 1.25	LR	LR	LR	0	00380038h	30112222h
3	800	1.5, 1.35, 1.25	NP	NP	LR	0	00390039h	10332222h
3	800	1.5, 1.35, 1.25	LR	NP	LR	0	003A003Ah	30222222h
3	800	1.5, 1.35, 1.25	LR	LR	LR	0	00360036h	30112222h
3	1066	1.5, 1.35, 1.25	NP	NP	LR	0	00373C37h	20332222h
3	1066	1.5, 1.35, 1.25	LR	NP	LR	0	00383C38h	30222222h
3	1066	1.5, 1.35	LR	LR	LR	0	00333C33h	30112222h
3	1333	1.5, 1.35, 1.25	NP	NP	LR	0	00353A35h	30332222h
3	1333	1.5, 1.35	LR	NP	LR	0	00363A36h	30222222h

Table 75: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (G34r1 & LRDIMM)

Condition						D18F2x94_dct[1:0] SlowAccessMode	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
3	1333	1.5	LR	LR	LR	0	00303A30h	30112222h
3	1600	1.5	NP	NP	LR	0	00333933h	30332222h
3	1600	1.5	LR	NP	LR	0	00343934h	30222222h

Table 76: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & UDIMM)

Condition						D18F2x94_dct[1:0] SlowAccessMode	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
1	667	1.5, 1.35, 1.25	SR	-	-	0	00000000h	00112222h
1	667	1.5, 1.35, 1.25	DR	-	-	0	003B0000h	00112222h
1	800	1.5, 1.35, 1.25	SR	-	-	0	00000000h	10112222h
1	800	1.5, 1.35, 1.25	DR	-	-	0	003B0000h	10112222h
1	1066	1.5, 1.35, 1.25	SR, DR	-	-	0	00383837h	20112222h
1	1333	1.5, 1.35, 1.25	SR, DR	-	-	0	00363635h	30112222h
1	1600	1.5	SR	-	-	0	00353533h	30112222h
1	1600	1.5	DR	-	-	1	00003533h	30112222h
2	667	1.5, 1.35, 1.25	NP	SR	-	0	00000000h	00112222h
2	667	1.5, 1.35, 1.25	NP	DR	-	0	003B0000h	00112222h
2	667	1.5, 1.35, 1.25	SR, DR	SR, DR	-	0	00390039h	10222222h
2	800	1.5, 1.35, 1.25	NP	SR	-	0	00000000h	10112222h
2	800	1.5, 1.35, 1.25	NP	DR	-	0	003B0000h	10112222h

Table 76: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & UDIMM)

Condition						D18F2x94_dctl[1:0]	D18F2x9C_x0000_0004_dctl[1:0]	D18F2x9C_x0000_0000_dctl[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
2	800	1.5, 1.35, 1.25	SR, DR	SR, DR	-	0	00390039h	20222222h
2	1066	1.5, 1.35, 1.25	NP	SR, DR	-	0	00383837h	20112222h
2	1066	1.5, 1.35, 1.25	SR, DR	SR, DR	-	0	003A3A3Ah	30222222h
2	1333	1.5, 1.35, 1.25	NP	SR, DR	-	0	00363635h	30112222h
2	1333	1.5, 1.35	SR	SR	-	1	00003939h	30222222h
2	1600	1.5	NP	SR	-	0	00353533h	30112222h
2	1600	1.5	NP	DR	-	1	00003533h	30112222h
3	667	1.5, 1.35, 1.25	NP	NP	SR	0	00000000h	00332222h
3	667	1.5, 1.35, 1.25	NP	NP	DR	0	003B0000h	00332222h
3	667	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00390039h	10222222h
3	800	1.5, 1.35, 1.25	NP	NP	SR	0	00000000h	10332222h
3	800	1.5, 1.35, 1.25	NP	NP	DR	0	003B0000h	10332222h
3	800	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00390039h	20222222h
3	1066	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00383837h	20332222h
3	1066	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	003A3A3Ah	30222222h
3	1333	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00363635h	30112222h
3	1333	1.5, 1.35	SR	NP	SR	1	00003939h	30222222h
3	1600	1.5	NP	NP	SR	0	00353533h	30332222h
3	1600	1.5	NP	NP	DR	1	00003533h	30332222h

Table 77: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & RDIMM)

Condition						D18F2x94_dct[1:0]	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
1	667	1.5, 1.35, 1.25	SR, DR	-	-	0	00000000h	00112222h
1	667	1.5, 1.35, 1.25	QR	-	-	0	00000000h	00222222h
1	800	1.5, 1.35, 1.25	SR, DR	-	-	0	00000000h	10112222h
1	800	1.5, 1.35, 1.25	QR	-	-	0	00000000h	10222222h
1	1066	1.5, 1.35, 1.25	SR, DR	-	-	0	003C3C3Ch	20112222h
1	1066	1.5, 1.35, 1.25	QR	-	-	0	003C3C3Ch	30222222h
1	1333	1.5, 1.35, 1.25	SR, DR	-	-	0	003A3A3Ah	30112222h
1	1333	1.5	QR	-	-	0	003A3A3Ah	30222222h
1	1600	1.5	SR, DR	-	-	0	00393939h	30112222h
2	667	1.5, 1.35, 1.25	NP	SR, DR	-	0	00000000h	00112222h
2	667	1.5, 1.35, 1.25	NP	QR	-	0	00000000h	00222222h
2	667	1.5, 1.35, 1.25	SR, DR, QR	SR, DR, QR	-	0	00000000h	10222222h
2	800	1.5, 1.35, 1.25	NP	SR, DR	-	0	00000000h	10112222h
2	800	1.5, 1.35, 1.25	NP	QR	-	0	00000000h	10222222h
2	800	1.5, 1.35, 1.25	SR, DR, QR	SR, DR, QR	-	0	00000000h	20222222h
2	1066	1.5, 1.35, 1.25	NP	SR, DR	-	0	00393C39h	20112222h
2	1066	1.5, 1.35, 1.25	NP	QR	-	0	00393C39h	20222222h
2	1066	1.5, 1.35, 1.25	SR, DR	SR, DR	-	0	003A3C3Ah	30222222h
2	1333	1.5, 1.35, 1.25	NP	SR, DR	-	0	00373A37h	30112222h
2	1333	1.5	NP	QR	-	0	00373A37h	30222222h
2	1333	1.5, 1.35	SR	SR	-	0	00383A38h	30222222h
2	1600	1.5	NP	SR, DR	-	0	00363936h	30112222h
3	667	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00000000h	00332222h
3	667	1.5, 1.35, 1.25	NP	QR	NP	0	00000000h	10222222h
3	667	1.5, 1.35, 1.25	NP	QR	SR, DR	0	00000000h	20222222h
3	667	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00000000h	10222222h

Table 77: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & RDIMM)

Condition						D18F2x94_dctl[1:0] SlowAccessMode	D18F2x9C_x0000_0004_dctl[1:0]	D18F2x9C_x0000_0000_dctl[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
3	667	1.5, 1.35, 1.25	SR, DR	SR, DR, QR	SR, DR	0	00380038h	30112222h
3	800	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00390039h	10332222h
3	800	1.5, 1.35, 1.25	NP	QR	NP	0	00390039h	20222222h
3	800	1.5, 1.35, 1.25	NP	QR	SR, DR	0	003A003Ah	30222222h
3	800	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	003A003Ah	20222222h
3	800	1.5, 1.35	SR, DR	SR, DR	SR, DR	0	00360036h	30112222h
3	800	1.25	SR	SR	SR	0	00360036h	30112222h
3	1066	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00373C37h	20332222h
3	1066	1.5	NP	QR	NP	0	00373C37h	30222222h
3	1066	1.5, 1.35, 1.25	SR, DR	NP	SR, DR	0	00383C38h	30222222h
3	1066	1.5	SR	SR	SR	0	00333C33h	30112222h
3	1333	1.5, 1.35, 1.25	NP	NP	SR, DR	0	00353A35h	30332222h
3	1333	1.5, 1.35	SR	NP	SR	0	00363A36h	30222222h
3	1600	1.5	NP	NP	SR, DR	0	00333933h	30332222h

Table 78: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (C32r1 & LRDIMM)

Condition						D18F2x94_dct[1:0]	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2			
1	667	1.5, 1.35, 1.25	LR	-	-	0	00000000h	00112222h
1	800	1.5, 1.35, 1.25	LR	-	-	0	00000000h	10112222h
1	1066	1.5, 1.35, 1.25	LR	-	-	0	003C3C3Ch	20112222h
1	1333	1.5, 1.35, 1.25	LR	-	-	0	003A3A3Ah	30112222h
1	1600	1.5, 1.35	LR	-	-	0	00393939h	30112222h
1	1866	1.5	LR	-	-	0	00393939h	30332222h
2	667	1.5, 1.35, 1.25	NP	LR	-	0	00000000h	00112222h
2	667	1.5, 1.35, 1.25	LR	LR	-	0	00000000h	10222222h
2	800	1.5, 1.35, 1.25	NP	LR	-	0	00000000h	10112222h
2	800	1.5, 1.35, 1.25	LR	LR	-	0	00000000h	20222222h
2	1066	1.5, 1.35, 1.25	NP	LR	-	0	00393C39h	20112222h
2	1066	1.5, 1.35, 1.25	LR	LR	-	0	003A3C3Ah	30222222h
2	1333	1.5, 1.35, 1.25	NP	LR	-	0	00373A37h	30112222h
2	1333	1.5, 1.35	LR	LR	-	0	00383A38h	30222222h
2	1600	1.5	NP	LR	-	0	00363936h	30112222h
3	667	1.5, 1.35, 1.25	NP	NP	LR	0	00000000h	00332222h
3	667	1.5, 1.35, 1.25	LR	NP	LR	0	00000000h	20222222h
3	667	1.5, 1.35, 1.25	LR	LR	LR	0	00380038h	30112222h
3	800	1.5, 1.35, 1.25	NP	NP	LR	0	00390039h	10332222h
3	800	1.5, 1.35, 1.25	LR	NP	LR	0	003A003Ah	30222222h
3	800	1.5, 1.35, 1.25	LR	LR	LR	0	00360036h	30112222h
3	1066	1.5, 1.35, 1.25	NP	NP	LR	0	00373C37h	20332222h
3	1066	1.5, 1.35, 1.25	LR	NP	LR	0	00383C38h	30222222h
3	1066	1.5	LR	LR	LR	0	00333C33h	30112222h
3	1333	1.5, 1.35, 1.25	NP	NP	LR	0	00353A35h	30332222h
3	1333	1.5, 1.35	LR	NP	LR	0	00363A36h	30222222h
3	1600	1.5	NP	NP	LR	0	00333933h	30332222h

Table 79: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (AM3r2 & UDIMM)

Condition					D18F2x94_dct[1:0]	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1			
1	667, 800	1.5, 1.35	SR	-	0	00000000h	00112222h
1	667, 800	1.5, 1.35	DR	-	0	003B0000h	00112222h
1	1066	1.5, 1.35	SR	-	0	00000000h	10112222h
1	1066	1.5, 1.35	DR	-	0	00380000h	10112222h
1	1333	1.5, 1.35	SR	-	0	00000000h	20112222h
1	1333	1.5, 1.35	DR	-	0	00360000h	20112222h
1	1600	1.5	SR	-	0	00000000h	30112222h
1	1600	1.5	DR	-	1	00000000h	30112222h
1	1866	1.5	SR	-	0	00000000h	30332222h
1	1866	1.5	DR	-	1	00000000h	30332222h
2	667, 800	1.5, 1.35	NP	SR	0	00000000h	00112222h
2	667, 800	1.5, 1.35	NP	DR	0	003B0000h	00112222h
2	667	1.5, 1.35	SR, DR	SR, DR	0	00390039h	10222322h
2	800	1.5, 1.35	SR, DR	SR, DR	0	00390039h	20222322h
2	1066	1.5, 1.35	NP	SR	0	00000000h	10112222h
2	1066	1.5, 1.35	NP	DR	0	00380000h	10112222h
2	1066	1.5, 1.35	SR, DR	SR, DR	0	00350037h	30222322h
2	1333	1.5, 1.35	NP	SR	0	00000000h	20112222h
2	1333	1.5, 1.35	NP	DR	0	00360000h	20112222h
2	1333	1.5	SR, DR	SR, DR	1	00000035h	30222322h
2	1333	1.35	SR	SR	1	00000035h	30222322h
2	1600	1.5	NP	SR	0	00000000h	30112222h
2	1600	1.5	NP	DR	1	00000000h	30112222h
2	1600	1.5	SR	SR	1	00000033h	30222322h

Table 80: BIOS Recommendations for SlowAccessMode, Addr/Cmd Timing (SODIMM)

Condition						D18F2x94_dct[1:0]	D18F2x9C_x0000_0004_dct[1:0]	D18F2x9C_x0000_0000_dct[1:0]
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	SlowAccessMode			
1	667, 800	1.5, 1.35	SR, DR	-	0	00000000h	00002222h	
1	1066	1.5, 1.35	SR	-	0	003D3D3Dh	10002222h	
1	1066	1.5, 1.35	DR	-	0	00000000h	10002222h	
1	1333	1.5, 1.35	SR	-	0	003D3D3Dh	20112222h	
1	1333	1.5, 1.35	DR	-	0	00003D3Dh	20112222h	
1	1600	1.5	SR	-	0	003C3C3Ch	30332222h	
1	1600	1.5	DR	-	1	00003C3Ch	30332222h	
1	1600	1.35	SR	-	0	003C3C3Ch	30332222h	
1	1600	1.35	DR	-	1	00003C3Ch	30332222h	
2	667, 800	1.5, 1.35	NP	SR, DR	0	00000000h	00002222h	
2	667	1.5, 1.35	SR, DR	SR, DR	1	00000039h	10222323h	
2	800	1.5, 1.35	SR, DR	SR, DR	1	00000039h	20222323h	
2	1066	1.5, 1.35	NP	SR	0	003D3D3Dh	10002222h	
2	1066	1.5, 1.35	NP	DR	0	00000000h	10002222h	
2	1066	1.5, 1.35	SR, DR	SR, DR	1	00000037h	30222323h	
2	1333	1.5, 1.35	NP	SR	0	003D3D3Dh	20112222h	
2	1333	1.5, 1.35	NP	DR	0	00003D3Dh	20112222h	
2	1333	1.5	SR, DR	SR, DR	1	00000035h	30222323h	
2	1333	1.35	SR	SR	1	00000035h	30222323h	

2.10.5.6 DCT Training Specific Configuration

The DCT requires certain features be disabled during DRAM device initialization and training. BIOS should program the registers in [Table 81](#) before DRAM device initialization and training. For normal operation, BIOS programs the recommended values if provided in [Table 81](#). BIOS must quiesce all other forms of DRAM traffic on the channel being trained. See [2.10.5 \[DCT/DRAM Initialization and Resume\]](#).

Table 81: 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	Fh
D18F2x94_dct[1:0][PowerDownEn]	0	1
D18F2x94_dct[1:0][ZqcsInterval]	00b	10b
D18F2x9C_x0000_000D_dct[1:0][RxMaxDurDII_NoLock]	000b	See 2.10.5.10
D18F2x9C_x0000_000D_dct[1:0][TxMaxDurDII_NoLock]	000b	See 2.10.5.10
D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0][EnRxPadStandby]	0	See 2.10.5.10
D18F2xA4[BwCapEn]	0	See 2.10.10
D18F2xA4[ODTSEn]	0	See 2.10.10
D18F2x110[DctSelIntLvEn]	0	x ¹
D18F3x58[L3Scrub]	0	See 2.13.1.8
D18F3x58[DramScrub]	0	See 2.13.1.8
D18F3x5C[ScrubReDirEn]	0	IF (D18F3x44[DramEccEn]==1) THEN 1 ELSE 0 ENDIF
D18F3x1B8[L3ScrbRedirDis]	1	0
1. Programmed specific to the current platform or memory configuration.		

2.10.5.7 DRAM Device and Controller Initialization

BIOS initializes the DRAM devices and the controller using a software controlled sequence. See 2.10.5.7.1 [Software DDR3 Device Initialization].

BIOS must observe additional requirements for changing the PLL frequency when setting D18F2x7C_dct[1:0][EnDramInit]. See 2.10.5.3.2 [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.

2.10.5.7.1 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. Program D18F2x7C_dct[1:0][EnDramInit] = 1.
2. Wait 200 us.
3. Program D18F2x7C_dct[1:0][DeassertMemRstX] = 1.
4. Wait 500 us.
5. Program D18F2x7C_dct[1:0][AssertCke] = 1.

6. Wait 360 ns.

The following step is performed with registered DIMMs or LRDIMMs:

7. Program `D18F2x90_dct[1:0][ParEn] = ~D18F2x90_dct[1:0][UnbuffDimm]`.

The following steps are performed with registered DIMMs only and must be performed for each chip select pair:

8. Send RC0, RC1, and RC2.
9. Wait 6 us.
10. Send RC3, RC4, and RC5.
11. Send RC6 and RC7 for custom settings at this time, as directed by the DIMM manufacturer's data sheet.
12. Send the remaining RCW. See 2.10.5.7.1.2.

The following steps are performed with LRDIMMs only and must be performed for each DIMM:

13. Send F0RC0, F0RC1, F0RC2, F0RC10, F0RC11, F1RC8, and F1RC11 - F1RC15.
14. Wait 6 us.
15. Send F0RC3, F0RC4, and F0RC5.
16. Send the remaining RCWs. See 2.10.5.7.1.2.
17. Send the extended RCWs. See 2.10.5.7.1.2.3

The following steps are performed once for each channel with unbuffered DIMMs, once for each chip select with registered DIMMs, and once for each physical rank with LRDIMMs:

18. Send MRS(2).
19. Send MRS(3). Ordinarily at this time, `MrsAddress[2:0] = 000b`.
20. Send MRS(1) with `MrsAddress[7] = 0`.
21. Send MRS(0) with `MrsAddress[8] = 1`.

The following steps are performed for all DIMM types:

22. 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.
23. Program `D18F2x7C_dct[1:0][EnDramInit] = 0`.

The following steps are performed with LRDIMMs only and must be performed for each DIMM:

24. Configure the DCT and LRDIMM:
 - Program `D18F2xA8_dct[1:0][LrDimmErrOutMonEn]=1`.
 - Send F2RC3 with data = 1000b.
25. Send F0RC12 with data = 0010b.
26. Wait until `D18F2xA0_dct[1:0][RcvParErr]=0` or the JEDEC specified training time expires.
27. Configure for normal operation:
 - Program `D18F2xA8_dct[1:0][LrDimmErrOutMonEn]=0`.
 - Send F2RC3 with data = 0000b.
 - Send F0RC12 with data = 0000b.

2.10.5.7.1.1 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 below:
 - MrsBank[2:0] = BA2:BA0 to address MR[3:0].
 - MrsAddress[15:0] = A15:A0 to write data[15:0].
 - BIOS may need to remap bits, see also [D18F2x\[5C:40\]_dct\[1:0\]\[OnDimmMirror\]](#).
 - Set all other bits in MrsAddress to zero. BIOS should write reserved fields as 0.
2. Program MrsChipSel as appropriate.
3. Program SendMrsCmd = 1.
4. Wait for SendMrsCmd = 0.

For LRDIMMs, the behavior of MRS commands issued to physical ranks associated with the logical rank selected by [D18F2x7C_dct\[1:0\]\[MrsChipSel\]](#) depends on the setting of [F0RC14\[MRSCCommandControl\]](#):

- MRSCCommandControl=0
 - The MRS command is issued to each physical rank of the LRDIMM associated with the logical rank.
- MRSCCommandControl=1
 - If MrsAddress[13] = 1
 - The MRS command is issued to each physical rank of the LRDIMM associated with the logical rank.
 - Else
 - The MRS command issued to a specific physical rank based on [D18F2x7C_dct\[1:0\]\[MrsAddress\[17:14\]\]](#). See [2.10.11.1 \[LRDIMM Rank Multiplication\]](#).

MR0_dct[1:0] DDR3 MR0

Table 82: BIOS Recommendations for MR0_dct[1:0][WR]

Condition	MR0_dct[1:0]
	WR
D18F2x22C_dct[1:0][Twr]	
10h	000b
5h	001b
6h	010b
7h	011b
8h	100b
Ah	101b
Ch	110b
Eh	111b

Table 83: BIOS Recommendations for MR0_dct[1:0][CL[3:0]]

Condition	MR0_dct[1:0]
	CL[3:0]
D18F2x200_dct[1:0][Tcl]	
5h	2h
6h	4h
7h	6h
8h	8h
9h	Ah

Table 83: BIOS Recommendations for MR0_dct[1:0][CL[3:0]]

Condition	MR0_dct[1:0]
D18F2x200_dct[1:0][Tcl]	CL[3:0]
Ah	Ch
Bh	Eh
Ch	1h
Dh	3h
Eh	5h
Fh	7h
10h	9h

Bits	Description
15:13	Reserved.
12	PPD: DLL control for precharge powerdown. BIOS: D18F2x84_dct[1:0][PchgPDMModeSel].
11:9	WR: write recovery for autoprecharge. BIOS: Table 82.
8	DLL: DLL reset. BIOS: See 2.10.5.7.1.
7	TM: test mode. BIOS: 0.
6:4	CL[3:1]. CAS latency. See: CL[0].
3	RBT: read burst type. BIOS: 1.
2	CL[0]. CAS latency. CL[3:0] = {CL[3:1], CL[0]}. BIOS: Table 83.
1:0	BL: burst length. BIOS: D18F2x84_dct[1:0][BurstCtrl].

MR1_dct[1:0] DDR3 MR1

Bits	Description
15:13	Reserved.
12	Qoff: Qoff. BIOS: See 2.10.5.8.1.
11	TDQS: TDQS enable. BIOS: IF (DeviceWidth == 001b & (mixed channel of x4 and x8 DIMMs) & ~LRDIMM) THEN 1 ELSE 0 ENDIF.
10	Reserved.
9	RttNom[2]: RttNom. See RttNom[0].
8	Reserved.
7	Level: write leveling enable. BIOS: See 2.10.5.8.1.
6	RttNom[1]: RttNom. See RttNom[0].
5	DIC[1]: output driver impedance control. See: DIC[0].
4:3	AL: additive latency. BIOS: 0.
2	RttNom[0]: RttNom. RttNom[2:0] = {RttNom[2], RttNom[1], RttNom[0]}. BIOS: IF LRDIMM THEN See 2.10.5.4 ELSE Table 56, Table 57, Table 59, Table 60, Table 62 ENDIF.

1	DIC[0]: output driver impedance control. DIC[1:0] = {DIC[1], DIC[0]}. BIOS: IF LRDIMM THEN See 2.10.5.4 ELSE 1 ENDIF.
0	DLL: DLL enable. BIOS: 0.

MR2_dct[1:0] DDR3 MR2

Table 84: BIOS Recommendations for MR2_dct[1:0][ASR, SRT]

Condition		MR2_dct[1:0]	
AutoSelfRefresh	ExtendedTemperature-Range	ASR	SRT
0	0	0	0
0	1	0	1
1	-	1	0

Bits	Description
15:11	Reserved.
10:9	RttWr: RttWr. BIOS: IF LRDIMM THEN See 2.10.5.4 ELSE Table 56 , Table 57 , Table 59 , Table 60 , Table 62 ENDIF.
8	Reserved.
7	SRT: self refresh temperature range. BIOS: Table 84 .
6	ASR: auto self refresh. BIOS: Table 84 .
5:3	CWL: CAS write latency. BIOS: D18F2x20C_dct[1:0][Tcwl] - 5.
2:0	PASR: partial array self refresh. BIOS: 0.

MR3 DDR3 MR3

Bits	Description
15:3	Reserved.
2	MPR: MPR operation. BIOS: 0.
1:0	MPRLoc: MPR location. BIOS: 0.

2.10.5.7.1.2 Software Control Word Initialization

DDR3 register devices on RDIMMs contain up to 16 control words, referred to as RC0 to RC15. LRDIMM buffer devices contain up to 16 functions, each containing up to 16 control words. Each control word is four bits. These devices are programmed at the bus using an RCW command by (a) presenting the 4-bit address of the control word on [BA2, A2, A1, A0], (b) presenting the 4-bit write data on [BA1, BA0, A4, A3], and (c) asserting both chip selects of a chip select pair. For LRDIMMs, RC7 serves as the function select control word.

A RCW(n) command for DDR3 register device initialization is accomplished by programming [D18F2x7C_dct\[1:0\]](#) and [D18F2xA8_dct\[1:0\]\[CtrlWordCS\]](#) as follows:

1. Program MrsBank and MrsAddress.

- n = [BA2, A2, A1, A0].
 - data = [BA1, BA0, A4, A3].
 - Set all other bits in MrsAddress to zero. BIOS should write reserved fields as 0.
2. Program `D18F2xA8_dct[1:0][CtrlWordCS]`=bit mask for the target chip selects.
 3. Set `SendControlWord` = 1.
 4. Wait for `SendControlWord` = 0.

For RDIMMs with `NumRegisters`=2, BIOS must initialize the register control words of both devices on the DIMM by programming each chip select pair.

BIOS must wait 6 us after programming `RC2`, `RC10`, `F0RC2`, `F0RC10`, `F0RC11`, `F1RC8`, or `F1RC11-F1RC15`.

2.10.5.7.1.2.1 RDIMM Control Words

RC0 DDR3 Register Control Word 0

Bits	Description
3:2	OutputsDisabled: outputs disabled. BIOS: 0.
1	FloatingOutputs: floating outputs. BIOS: 1.
0	OutputInversion: output inversion. BIOS: 0.

RC1 DDR3 Register Control Word 1

Bits	Description
3:0	DisableOutputClock: disable output clock. BIOS: IF (DIMM == SR) THEN Ch ELSE 0 ENDIF.

RC2 DDR3 Register Control Word 2

Table 85: BIOS Recommendations for `RC2[IBT]` (G34r1)

Condition								RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	NumRegisters	IBT[3:0]
1	667, 800	1.5, 1.35, 1.25	SR	-	-	SR	1	1h
1	667, 800	1.5, 1.35, 1.25	DR	-	-	DR	1	1h
1	667, 800	1.5, 1.35, 1.25	QR	-	-	QR	1, 2	1h
1	1066, 1333	1.5, 1.35, 1.25	SR	-	-	SR	1	0h
1	1066, 1333	1.5, 1.35, 1.25	DR	-	-	DR	1	0h
1	1066	1.5, 1.35, 1.25	QR	-	-	QR	1	0h
1	1066	1.5, 1.35, 1.25	QR	-	-	QR	2	1h

Table 85: BIOS Recommendations for RC2[IBT] (G34r1)

Condition								RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	NumRegisters	IBT[3:0]
1	1333	1.5, 1.35	QR	-	-	QR	1	0h
1	1333	1.5, 1.35	QR	-	-	QR	2	1h
1	1600	1.5, 1.35	SR	-	-	SR	1	0h
1	1600	1.5, 1.35	DR	-	-	DR	1	0h
1	1866	1.5	SR	-	-	SR	1	0h
1	1866	1.5	DR	-	-	DR	1	0h
2	667, 800	1.5, 1.35, 1.25	NP, SR	SR	-	SR	1	1h
2	667, 800	1.5, 1.35, 1.25	NP, DR	DR	-	DR	1	1h
2	667, 800	1.5, 1.35, 1.25	NP	QR	-	QR	1, 2	1h
2	667, 800, 1066	1.5, 1.35, 1.25	SR	DR	-	SR, DR	1	1h
2	667, 800	1.5, 1.35, 1.25	SR	QR	-	SR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	SR, DR	QR	-	QR	2	8h
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR	-	SR, DR	1	1h
2	667, 800	1.5, 1.35, 1.25	DR	QR	-	DR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	QR	SR	-	SR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	QR	SR, DR, QR	-	QR	2	8h
2	667, 800	1.5, 1.35, 1.25	QR	DR	-	DR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	QR	QR	-	QR	1	1h
2	1066, 1333	1.5, 1.35, 1.25	NP	SR	-	SR	1	0h
2	1066, 1333	1.5, 1.35, 1.25	NP	DR	-	DR	1	0h
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	1	0h
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	2	1h
2	1066, 1333	1.5, 1.35, 1.25	SR	SR	-	SR	1	1h
2	1066	1.5	SR	QR	-	SR, QR	1	1h
2	1066	1.5	SR, DR	QR	-	QR	2	8h
2	1066	1.5, 1.35, 1.25	DR	DR	-	DR	1	1h
2	1066	1.5	DR	QR	-	DR, QR	1	1h
2	1066	1.5	QR	SR	-	SR, QR	1	1h
2	1066	1.5	QR	SR, DR, QR	-	QR	2	8h
2	1066	1.5	QR	DR	-	DR, QR	1	1h
2	1066	1.5	QR	QR	-	QR	1	1h

Table 85: BIOS Recommendations for RC2[IBT] (G34r1)

Condition								RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	NumRegisters	IBT[3:0]
2	1333	1.5	NP	QR	-	QR	1	0h
2	1333	1.5	NP	QR	-	QR	2	1h
2	1333	1.5, 1.35	SR	DR	-	SR, DR	1	1h
2	1333	1.5, 1.35	DR	SR	-	SR, DR	1	1h
2	1333	1.5, 1.35	DR	DR	-	DR	1	1h
2	1600	1.5	NP	SR	-	SR	1	0h
2	1600	1.5	NP	DR	-	DR	1	0h
2	1600	1.5	SR	SR	-	SR	1	1h
2	1600	1.5	SR	DR	-	SR, DR	1	1h
2	1600	1.5	DR	SR	-	SR, DR	1	1h
2	1600	1.5	DR	DR	-	DR	1	1h
3	667, 800	1.5, 1.35, 1.25	NP	NP	SR	SR	1	1h
3	667	1.5, 1.35, 1.25	NP	NP	DR	DR	1	1h
3	667, 800	1.5, 1.35, 1.25	NP	QR	NP	QR	1, 2	1h
3	667	1.5, 1.35, 1.25	NP, SR	QR	SR	SR, QR	1	1h
3	667	1.5, 1.35, 1.25	NP, SR, DR	QR	SR, DR	QR	2	8h
3	667	1.5, 1.35, 1.25	NP, DR	QR	DR	DR, QR	1	1h
3	667, 800	1.5, 1.35, 1.25	SR	NP, SR	SR	SR	1	1h
3	667	1.5, 1.35, 1.25	SR	NP, SR	DR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	SR	DR	SR, DR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	SR	QR	DR	SR, DR, QR	1	1h
3	667	1.5, 1.35, 1.25	DR	NP, DR	SR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	DR	NP, DR	DR	DR	1	1h
3	667	1.5, 1.35, 1.25	DR	SR	SR, DR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	DR	QR	SR	SR, DR, QR	1	1h
3	800	1.5, 1.35	NP	NP	DR	DR	1	1h
3	800	1.5, 1.35	NP, SR	QR	SR	SR, QR	1	1h
3	800	1.5, 1.35	NP, SR, DR	QR	SR, DR	QR	2	8h
3	800	1.5, 1.35	NP, DR	QR	DR	DR, QR	1	1h
3	800	1.5, 1.35	SR	NP, SR	DR	SR, DR	1	1h
3	800	1.5, 1.35	SR	DR	SR, DR	SR, DR	1	1h

Table 85: BIOS Recommendations for RC2[IBT] (G34r1)

Condition								RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	NumRegisters	IBT[3:0]
3	800	1.5, 1.35	SR	QR	DR	SR, DR, QR	1	1h
3	800	1.5, 1.35	DR	NP, DR	SR	SR, DR	1	1h
3	800	1.5, 1.35	DR	NP, DR	DR	DR	1	1h
3	800	1.5, 1.35	DR	SR	SR, DR	SR, DR	1	1h
3	800	1.5, 1.35	DR	QR	SR	SR, DR, QR	1	1h
3	800	1.25	NP	QR	SR	SR, QR	1	1h
3	800	1.25	NP	QR	SR, DR	QR	2	8h
3	800	1.25	NP	QR	DR	DR, QR	1	1h
3	800	1.25	NP, DR	NP	DR	DR	1	1h
3	800	1.25	SR	NP	DR	SR, DR	1	1h
3	800	1.25	DR	NP	SR	SR, DR	1	1h
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	SR	SR	1	0h
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	DR	DR	1	0h
3	1066	1.5, 1.35	NP	QR	NP	QR	1	0h
3	1066	1.5, 1.35	NP	QR	NP	QR	2	1h
3	1066	1.5, 1.35	SR	NP, SR	SR	SR	1	1h
3	1066	1.5	SR	NP, SR	DR	SR, DR	1	1h
3	1066	1.5	SR	DR	SR, DR	SR, DR	1	1h
3	1066	1.5	DR	NP, DR	SR	SR, DR	1	1h
3	1066	1.5	DR	NP, DR	DR	DR	1	1h
3	1066	1.5	DR	SR	SR, DR	SR, DR	1	1h
3	1066	1.35, 1.25	SR	NP	DR	SR, DR	1	1h
3	1066	1.35, 1.25	DR	NP	SR	SR, DR	1	1h
3	1066	1.35, 1.25	DR	NP	DR	DR	1	1h
3	1066	1.25	SR	NP	SR	SR	1	1h
3	1333	1.5, 1.35, 1.25	SR	NP	SR	SR	1	1h
3	1333	1.5, 1.35	SR	NP	DR	SR, DR	1	1h
3	1333	1.5, 1.35	DR	NP	SR	SR, DR	1	1h
3	1333	1.5, 1.35	DR	NP	DR	DR	1	1h
3	1600	1.5	NP	NP	SR	SR	1	0h
3	1600	1.5	NP	NP	DR	DR	1	0h

Table 86: BIOS Recommendations for RC2[IBT] (C32r1)

Condition								RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	NumRegisters	IBT[3:0]
1	667, 800	1.5, 1.35, 1.25	SR	-	-	SR	1	1h
1	667, 800	1.5, 1.35, 1.25	DR	-	-	DR	1	1h
1	667, 800	1.5, 1.35, 1.25	QR	-	-	QR	1, 2	1h
1	1066, 1333	1.5, 1.35, 1.25	SR	-	-	SR	1	0h
1	1066, 1333	1.5, 1.35, 1.25	DR	-	-	DR	1	0h
1	1066	1.5, 1.35, 1.25	QR	-	-	QR	1	0h
1	1066	1.5, 1.35, 1.25	QR	-	-	QR	2	1h
1	1333	1.5	QR	-	-	QR	1	0h
1	1333	1.5	QR	-	-	QR	2	1h
1	1600	1.5	SR	-	-	SR	1	0h
1	1600	1.5	DR	-	-	DR	1	0h
2	667, 800	1.5, 1.35, 1.25	NP, SR	SR	-	SR	1	1h
2	667, 800	1.5, 1.35, 1.25	NP, DR	DR	-	DR	1	1h
2	667, 800	1.5, 1.35, 1.25	NP	QR	-	QR	1, 2	1h
2	667, 800, 1066	1.5, 1.35, 1.25	SR	DR	-	SR, DR	1	1h
2	667, 800	1.5, 1.35, 1.25	SR	QR	-	SR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	SR, DR	QR	-	QR	2	8h
2	667, 800, 1066	1.5, 1.35, 1.25	DR	SR	-	SR, DR	1	1h
2	667, 800	1.5, 1.35, 1.25	DR	QR	-	DR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	QR	SR	-	SR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	QR	SR, DR, QR	-	QR	2	8h
2	667, 800	1.5, 1.35, 1.25	QR	DR	-	DR, QR	1	1h
2	667, 800	1.5, 1.35, 1.25	QR	QR	-	QR	1	1h
2	1066, 1333	1.5, 1.35, 1.25	NP	SR	-	SR	1	0h
2	1066, 1333	1.5, 1.35, 1.25	NP	DR	-	DR	1	0h
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	1	0h
2	1066	1.5, 1.35, 1.25	NP	QR	-	QR	2	1h
2	1066	1.5, 1.35, 1.25	SR	SR	-	SR	1	1h
2	1066	1.5, 1.35, 1.25	DR	DR	-	DR	1	1h
2	1333	1.5	NP	QR	-	QR	1	0h
2	1333	1.5	NP	QR	-	QR	2	1h
2	1333	1.5, 1.35	SR	SR	-	SR	1	1h

Table 86: BIOS Recommendations for RC2[IBT] (C32r1)

Condition								RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	NumRegisters	IBT[3:0]
2	1600	1.5	NP	SR	-	SR	1	0h
2	1600	1.5	NP	DR	-	DR	1	0h
3	667, 800	1.5, 1.35, 1.25	NP	NP	SR	SR	1	1h
3	667	1.5, 1.35, 1.25	NP	NP	DR	DR	1	1h
3	667, 800	1.5, 1.35, 1.25	NP	QR	NP	QR	1, 2	1h
3	667	1.5, 1.35, 1.25	NP, SR	QR	SR	SR, QR	1	1h
3	667	1.5, 1.35, 1.25	NP, SR, DR	QR	SR, DR	QR	2	8h
3	667	1.5, 1.35, 1.25	NP, DR	QR	DR	DR, QR	1	1h
3	667, 800	1.5, 1.35, 1.25	SR	NP, SR	SR	SR	1	1h
3	667	1.5, 1.35, 1.25	SR	NP, SR	DR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	SR	DR	SR, DR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	SR	QR	DR	SR, DR, QR	1	1h
3	667	1.5, 1.35, 1.25	DR	NP, DR	SR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	DR	NP, DR	DR	DR	1	1h
3	667	1.5, 1.35, 1.25	DR	SR	SR, DR	SR, DR	1	1h
3	667	1.5, 1.35, 1.25	DR	QR	SR	SR, DR, QR	1	1h
3	800	1.5, 1.35	NP	NP	DR	DR	1	1h
3	800	1.5, 1.35, 1.25	NP	QR	SR	SR, QR	1	1h
3	800	1.5, 1.35, 1.25	NP	QR	SR, DR	QR	2	8h
3	800	1.5, 1.35, 1.25	NP	QR	DR	DR, QR	1	1h
3	800	1.5, 1.35	SR	NP, SR	DR	SR, DR	1	1h
3	800	1.5, 1.35	SR	DR	SR, DR	SR, DR	1	1h
3	800	1.5, 1.35	DR	NP, DR	SR	SR, DR	1	1h
3	800	1.5, 1.35	DR	NP, DR	DR	DR	1	1h
3	800	1.5, 1.35	DR	SR	SR, DR	SR, DR	1	1h
3	800	1.25	NP, DR	NP	DR	DR	1	1h
3	800	1.25	SR	NP	DR	SR, DR	1	1h
3	800	1.25	DR	NP	SR	SR, DR	1	1h
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	SR	SR	1	0h
3	1066, 1333	1.5, 1.35, 1.25	NP	NP	DR	DR	1	0h
3	1066	1.5	NP	QR	NP	QR	1	0h
3	1066	1.5	NP	QR	NP	QR	2	1h

Table 86: BIOS Recommendations for RC2[IBT] (C32r1)

Condition								RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	DIMM	NumRegisters	IBT[3:0]
3	1066	1.5	SR	NP, SR	SR	SR	1	1h
3	1066	1.5, 1.35, 1.25	SR	NP	DR	SR, DR	1	1h
3	1066	1.5, 1.35, 1.25	DR	NP	SR	SR, DR	1	1h
3	1066	1.5, 1.35, 1.25	DR	NP	DR	DR	1	1h
3	1066	1.35, 1.25	SR	NP	SR	SR	1	1h
3	1333	1.5, 1.35	SR	NP	SR	SR	1	1h
3	1600	1.5	NP	NP	SR	SR	1	0h
3	1600	1.5	NP	NP	DR	DR	1	0h

Bits	Description
3	FrequencyBandSelect: frequency band select. BIOS: 0.
2	IBT[0]: input bus termination. IBT[3:0] = {RC8[IBT[3:1]], IBT[0]}. BIOS: Table 85 , Table 86 . BIOS settings of this field are for motherboards which meet the relevant motherboard design guidelines.
1	OutputTiming: output timing. BIOS: 0.
0	AddrCmdPrelaunch: address and command prelaunch. BIOS: 0.

RC3 DDR3 Register Control Word 3

Bits	Description
3:0	CommandAddressDriverOutputs: command and address driver outputs. BIOS: See 2.10.5.4 .

RC4 DDR3 Register Control Word 4

Bits	Description
3:0	ControlDriverOutputs: control driver outputs. BIOS: See 2.10.5.4 .

RC5 DDR3 Register Control Word 5

Bits	Description
3:0	ClockOutputDrivers: clock output drivers. BIOS: See 2.10.5.4 .

RC[7:6] DDR3 Register Control Word [7:6]

Bits	Description
3:0	Reserved.

RC8 DDR3 Register Control Word 8

Bits	Description
3	MirrorMode: mirror mode. BIOS: 0.
2:0	IBT[3:1]: input bus termination [3:1]. See: RC2[IBT[0]] .

RC9 DDR3 Register Control Word 9

Bits	Description
3	CKEPowerDownModeEnable: CKE power down mode enable. BIOS: 1.
2	CKEPowerDownMode: CKE power down mode. BIOS: IF (mixed SR and (DR or QR) DIMMs on the channel and this DIMM is SR) THEN 0 ELSE 1 ENDIF.
1	Reserved.
0	WeakDriveMode: weak drive mode. BIOS: 1.

RC10 DDR3 Register Control Word 10

Table 87: BIOS Recommendations for RC10[OperatingSpeed]

Condition	RC10
DdrRate	OperatingSpeed
667, 800	0h
1066	1h
1333	2h
1600	3h
1866	4h

Bits	Description
3	Reserved.
2:0	OperatingSpeed: operating speed. BIOS: See Table 87 .

RC11 DDR3 Register Control Word 11

Bits	Description
------	-------------

3:2	Reserved.
1:0	OperatingVoltage: operating voltage. BIOS: IF (DdrVDDIO == 1.5) THEN 00b ELSEIF (DdrVDDIO == 1.35) THEN 01b ELSE 10b ENDIF.

RC[15:12] DDR3 Register Control Word [15:12]

Bits	Description
3:0	Reserved.

2.10.5.7.1.2.2 LRDIMM Control Words

F0RC0 DDR3 Buffer Function 0 Control Word 0

Bits	Description
3:2	QVrefOutput: QVref output. BIOS: 0.
1	OutputWeakDrive: output weak drive. BIOS: 1.
0	OutputInversion: output inversion. BIOS: 0.

F0RC1 DDR3 Buffer Function 0 Control Word 1

Bits	Description
3:0	DisableOutputClock: disable output clock. BIOS: 0.

F0RC2 DDR3 Buffer Function 0 Control Word 2

Bits	Description
3	FrequencyBandSelect: frequency band select. BIOS: 0.
2	Reserved.
1	RankSwap: rank swap. BIOS: See 2.10.5.4 .
0	AddrCmdPrelaunch: address and command prelaunch. BIOS: See 2.10.5.4 .

F0RC3 DDR3 Buffer Function 0 Control Word 3

Bits	Description
3:2	CSDriverOutputs: CS driver outputs. BIOS: See 2.10.5.4 .
1:0	CommandAddressDriverOutputs: command and address driver outputs. BIOS: See 2.10.5.4 .

F0RC4 DDR3 Buffer Function 0 Control Word 4

Bits	Description
3:2	CKEDriverOutputs: CKE driver outputs. BIOS: See 2.10.5.4.
1:0	ODTDriverOutputs: ODT driver outputs. BIOS: See 2.10.5.4.

F0RC5 DDR3 Buffer Function 0 Control Word 5

Bits	Description
3:0	ClockOutputDrivers: clock output drivers. BIOS: See 2.10.5.4.

F0RC6 DDR3 Buffer Function 0 Control Word 6

Bits	Description
3	Reserved.
2	ODTControl: ODT control. BIOS: 0.
1:0	CKEControl: CKE control. BIOS: 0h.

F[15:0]RC7 DDR3 Buffer Function [15:0] Control Word 7

Bits	Description				
3:0	FunctionSelect: function select. <table border="0"> <tr> <td><u>Bits</u></td> <td><u>Description</u></td> </tr> <tr> <td>Fh-0h</td> <td>Select access to Function <FunctionSelect> RC[15:0]</td> </tr> </table>	<u>Bits</u>	<u>Description</u>	Fh-0h	Select access to Function <FunctionSelect> RC[15:0]
<u>Bits</u>	<u>Description</u>				
Fh-0h	Select access to Function <FunctionSelect> RC[15:0]				

F0RC8 DDR3 Buffer Function 0 Control Word 8

Table 88: BIOS Recommendations for IBT (G34r1)

Condition						F0RC8	F1RC0	F1RC1	F1RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	IBT	IBT	IBT	IBT
1	667, 800	1.5, 1.35, 1.25	LR	-	-	001b	001b	001b	001b
1	1066, 1333	1.5, 1.35, 1.25	LR	-	-	000b	000b	000b	000b
1	1600	1.5, 1.35	LR	-	-	000b	000b	000b	000b
1	1866	1.5	LR	-	-	000b	000b	000b	000b

Table 88: BIOS Recommendations for IBT (G34r1)

Condition						F0RC8	F1RC0	F1RC1	F1RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	IBT	IBT	IBT	IBT
2	667, 800	1.5, 1.35, 1.25	NP	LR	-	001b	001b	001b	001b
2	667, 800	1.5, 1.35, 1.25	LR	LR	-	000b	000b	000b	000b
2	1066	1.5, 1.35, 1.25	NP, LR	LR	-	000b	000b	000b	000b
2	1333	1.5, 1.35	NP, LR	LR	-	000b	000b	000b	000b
2	1333	1.25	NP	LR	-	000b	000b	000b	000b
2	1600	1.5	NP, LR	LR	-	000b	000b	000b	000b
3	667	1.5, 1.35, 1.25	NP, LR	NP	LR	001b	001b	001b	001b
3	667, 800	1.5, 1.35, 1.25	LR	LR	LR	000b	000b	000b	000b
3	800, 1066, 1333	1.5, 1.35, 1.25	NP	NP	LR	000b	000b	000b	000b
3	800, 1066	1.5, 1.35, 1.25	LR	NP	LR	001b	001b	001b	001b
3	1066	1.5, 1.35	LR	LR	LR	000b	000b	000b	000b
3	1333	1.5, 1.35	LR	NP	LR	001b	001b	001b	001b
3	1333	1.5	LR	LR	LR	000b	000b	000b	000b
3	1600	1.5	NP	NP	LR	000b	000b	000b	000b
3	1600	1.5	LR	NP	LR	001b	001b	001b	001b

Table 89: BIOS Recommendations for IBT (C32r1)

Condition						F0RC8	F1RC0	F1RC1	F1RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	IBT	IBT	IBT	IBT
1	667, 800	1.5, 1.35, 1.25	LR	-	-	001b	001b	001b	001b
1	1066, 1333	1.5, 1.35, 1.25	LR	-	-	000b	000b	000b	000b
1	1600	1.5, 1.35	LR	-	-	000b	000b	000b	000b
1	1866	1.5	LR	-	-	000b	000b	000b	000b
2	667, 800	1.5, 1.35, 1.25	NP	LR	-	001b	001b	001b	001b
2	667, 800	1.5, 1.35, 1.25	LR	LR	-	000b	000b	000b	000b
2	1066	1.5, 1.35, 1.25	NP, LR	LR	-	000b	000b	000b	000b
2	1333	1.5, 1.35	NP, LR	LR	-	000b	000b	000b	000b

Table 89: BIOS Recommendations for IBT (C32r1)

Condition						F0RC8	F1RC0	F1RC1	F1RC2
NumDimmSlots	DdrRate	DdrVDDIO	DIMM0	DIMM1	DIMM2	IBT	IBT	IBT	IBT
	2	1333	1.25	NP	LR	-	000b	000b	000b
2	1600	1.5	NP	LR	-	000b	000b	000b	000b
3	667	1.5, 1.35, 1.25	NP, LR	NP	LR	001b	001b	001b	001b
3	667, 800	1.5, 1.35, 1.25	LR	LR	LR	000b	000b	000b	000b
3	800, 1066, 1333	1.5, 1.35, 1.25	NP	NP	LR	000b	000b	000b	000b
3	800, 1066	1.5, 1.35, 1.25	LR	NP	LR	001b	001b	001b	001b
3	1066	1.5	LR	LR	LR	000b	000b	000b	000b
3	1333	1.5, 1.35	LR	NP	LR	001b	001b	001b	001b
3	1600	1.5	NP	NP	LR	000b	000b	000b	000b

Bits	Description
3	VrefCA: QVref command address. BIOS: 0.
2:0	IBT: input bus termination. BIOS: Table 88 , Table 89 . BIOS settings of this field are for motherboards which meet the relevant motherboard design guidelines.

F0RC9 DDR3 Buffer Function 0 Control Word 9

Bits	Description
3	CKEPowerDownModeEnable: CKE power down mode enable. BIOS: 1.
2	CKEPowerDownMode: CKE power down mode. BIOS: 1.
1	OutputClockCKEPowerDownMode: output clock disable in CKE power down mode. BIOS: 0.
0	Reserved.

F0RC10 DDR3 Buffer Function 0 Control Word 10
Table 90: BIOS Recommendations for F0RC10[OperatingSpeed]

Condition	F0RC10
DdrRate	OperatingSpeed
667, 800	0h
1066	1h

Table 90: BIOS Recommendations for F0RC10[OperatingSpeed]

Condition	F0RC10
DdrRate	OperatingSpeed
1333	2h
1600	3h
1866	4h

Bits	Description
3	Reserved.
2:0	OperatingSpeed: operating speed. BIOS: Table 90.

F0RC11 DDR3 Buffer Function 0 Control Word 11

Table 91: BIOS Recommendations for F0RC11[ParityCalculation]

Condition	F0RC11
F0RC15[RankMultiplicationControl]	ParityCalculation
3h, 6h, 7h	10b
0h, 1h, 2h, 5h	01b

Bits	Description
3:2	ParityCalculation: parity calculation. BIOS: Table 91.
1:0	OperatingVoltage: operating voltage. BIOS: IF(DdrVDDIO == 1.5) THEN 00b ELSEIF (DdrVD-DIO == 1.35) THEN 01b ELSE 10b ENDIF.

F0RC12 DDR3 Buffer Function 0 Control Word 12

Bits	Description										
3	ContextOperationTraining: context for operation and training. BIOS: 0.										
2:0	TrainingControl: training control. BIOS: See 2.10.5.7.1, 2.10.5.8.1. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Normal operating mode</td> </tr> <tr> <td>1h</td> <td>Host to Buffer write leveling mode</td> </tr> <tr> <td>2h</td> <td>Buffer to DRAM training</td> </tr> <tr> <td>7h-3h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	Normal operating mode	1h	Host to Buffer write leveling mode	2h	Buffer to DRAM training	7h-3h	Reserved
Bits	Description										
0h	Normal operating mode										
1h	Host to Buffer write leveling mode										
2h	Buffer to DRAM training										
7h-3h	Reserved										

F0RC13 DDR3 Buffer Function 0 Control Word 13

Table 92: BIOS Recommendations for F0RC13[NumPhysicalRanks]

Condition	F0RC13
NumRanks	NumPhysicalRanks
000b	11b
001b	10b
011b	01b
100b	00b

Table 93: BIOS Recommendations for F0RC13[NumLogicalRanks]

Condition	F0RC13		
F0RC13[NumPhysical-Ranks]	DramCapacity	NumDimmSlots	NumLogicalRanks
11b	-	-	00b
10b	-	-	01b
01b	-	1, 2	10b
	-	3	01b
00b	0100b	-	01b
	-	1, 2	10b
	-	3	01b

Bits	Description
3:2	NumLogicalRanks: number logical ranks. BIOS: Table 93 .
1:0	NumPhysicalRanks: number physical ranks. BIOS: Table 92 .

F0RC14 DDR3 Buffer Function 0 Control Word 14

Bits	Description
3	DramBusWidth: DRAM bus width. BIOS: IF (DeviceWidth==0) THEN 0 ELSE 1 ENDIF.
2	MRSCCommandControl: MRS command control. BIOS: IF (F0RC15[RankMultiplicationControl] > 0) THEN 1 ELSE 0 ENDIF.
1	RefreshPrechargeCommandControl: refresh and precharge command control. BIOS: IF (F0RC15[RankMultiplicationControl] > 0) THEN D18F2xA8_dct[1:0][LrDimmEnhRefEn] ELSE 0 ENDIF.
0	AddressMirror: address mirror. BIOS: RankMap . See D18F2x[5C:40]_dct[1:0][OnDimmMirror].

F0RC15 DDR3 Buffer Function 0 Control Word 15

Table 94: BIOS Recommendations for F0RC15[RankMultiplicationControl]

Condition	DramCapacity	F0RC15 RankMultiplicationControl
$((1 \ll \text{POW}(\text{F0RC13}[\text{NumPhysicalRanks}], 3\text{h}) / (1 \ll \text{F0RC13}[\text{NumLogicalRanks}])))$		
1	-	0h
2	2h	1h
	3h	2h
	4h	3h
4	2h	5h
	3h	6h
	4h	7h

Bits	Description
3:0	RankMultiplicationControl: rank multiplication control. BIOS: Table 94 .

F1RC0 DDR3 Buffer Function 1 Control Word 0

Bits	Description
3	IBTCS32Select: input bus termination CS[3:2] select. BIOS: D18F2xA8_dct[1:0][CsTimingMux67] .
2:0	IBT: input bus termination. BIOS: Table 88 , Table 89 . Controls IBT for the CS pins.

F1RC1 DDR3 Buffer Function 1 Control Word 1

Bits	Description
3	Reserved.
2:0	IBT: input bus termination. BIOS: Table 88 , Table 89 . Controls IBT for the CKE pins.

F1RC2 DDR3 Buffer Function 1 Control Word 2

Bits	Description
3	Reserved.
2:0	IBT: input bus termination. BIOS: Table 88 , Table 89 . Controls IBT for the ODT pins.

F1RC3 DDR3 Buffer Function 1 Control Word 3

Bits	Description
3:2	Reserved.
1:0	WeakOutputDrive: weak output drive. BIOS: 0.

F1RC[6:4] DDR3 Buffer Function 1 Control Word [6:4]

Bits	Description
3:0	Reserved.

F1RC8 DDR3 Buffer Function 1 Control Word 8

Bits	Description
3:2	TotalQCSDelay: total QCS delay. BIOS: See 2.10.5.4.
1:0	TotalYDelay: total Y delay. BIOS: See 2.10.5.4.

F1RC9 DDR3 Buffer Function 1 Control Word 9

Bits	Description
3:0	RefStagger: refresh stagger. BIOS: IF (D18F2xA8_dct[1:0][LrDimmEnhRefEn]) THEN 0 ELSE 1h ENDIF.

F1RC10 DDR3 Buffer Function 1 Control Word 10

Bits	Description
3	StartRank: start rank. BIOS: 0.
2:0	RefStaggerLimit: refresh stagger limit. BIOS: 0.

F1RC11 DDR3 Buffer Function 1 Control Word 11

Bits	Description
3:2	TotalQCKEDelay: total QCKE delay. BIOS: See 2.10.5.4.
1:0	TotalQODTDelay: total QODT delay. BIOS: See 2.10.5.4.

F1RC12 DDR3 Buffer Function 1 Control Word 12

Bits	Description
3	Reserved.
2:0	QCAPrelaunchDelay: QCA prelaunch delay. BIOS: See 2.10.5.4.

F1RC13 DDR3 Buffer Function 1 Control Word 13

Bits	Description
3	QCSDelayControl: QCS delay control. BIOS: See 2.10.5.4.
2:0	QCSDelay: QCS delay. BIOS: See 2.10.5.4.

F1RC14 DDR3 Buffer Function 1 Control Word 14

Bits	Description
3	QODTDelayControl: QODT delay control. BIOS: See 2.10.5.4.
2:0	QODTDelay: QODT delay. BIOS: See 2.10.5.4.

F1RC15 DDR3 Buffer Function 1 Control Word 15

Bits	Description
3	QCKEDelayControl: QCKE delay control. BIOS: See 2.10.5.4.
2:0	QCKEDelay: QCKE delay. BIOS: See 2.10.5.4.

F2RC0 DDR3 Buffer Function 2 Control Word 0

Bits	Description
3:2	CSMultiplicationTransparentMode: chip select multiplication Transparent Mode. BIOS: 0.
1	Reserved.
0	TransparentModeEn: transparent mode enable. BIOS: 0.

F2RC1 DDR3 Buffer Function 2 Control Word 1

Bits	Description
3	Reserved.
2	MaskDDRReset: mask DDR reset. BIOS: 0.
1	ClearStickyRegisterBits: clear sticky register bits. BIOS: 0.
0	SoftReset: soft reset. BIOS: 0.

F2RC2 DDR3 Buffer Function 2 Control Word 2

Bits	Description
3:0	SMBusAccessControl: SMBus access control. BIOS: Fh.

F2RC3 DDR3 Buffer Function 2 Control Word 3

Bits	Description
3	ErroutEnable: errout enable. BIOS: 0. See also D18F2xA8_dct[1:0][LrDimmErrOutMonEn] .
2:1	Reserved.
0	TrainingControl: training control. BIOS: 0.

F2RC[15:8,6:4] DDR3 Buffer Function 2 Control Word [15:8,6:4]

Bits	Description
3:0	Reserved.

F3RC0 DDR3 Buffer Function 3 Control Word 0

Bits	Description
3	TDQSControl: TDQS control. BIOS: 0.
2:0	RttNom: RttNom. BIOS: Table 58 , Table 61 .

F3RC1 DDR3 Buffer Function 3 Control Word 1

Bits	Description
3	Vref: Vref. BIOS: 0.
2:0	RttWr: RttWr. BIOS: Table 58 , Table 61 .

F3RC2 DDR3 Buffer Function 3 Control Word 2

Bits	Description
3	DriverDisable: driver disable. BIOS: 0.
2:0	DIC: driver impedance control. BIOS: 1.

F3RC[5:3] DDR3 Buffer Function 3 Control Word [5:3]

Bits	Description
3:0	Reserved.

F3RC6 DDR3 Buffer Function 3 Control Word 6

Bits	Description
3	DRAMWidth: DRAM width. BIOS: IF (D18F2x90_dct[1:0][X4Dimm]==0) THEN 1 ELSE 0 ENDIF.
2:1	Reserved.
0	DQTimingMode: DQ timing mode. BIOS: 0.

F3RC8 DDR3 Buffer Function 3 Control Word 8

Bits	Description
3	Vref: Vref. BIOS: 0.
2:0	ODTStrength: ODT strength. BIOS: See 2.10.5.4.

F3RC9 DDR3 Buffer Function 3 Control Word 9

Bits	Description
3	DriverDisable: driver disable. BIOS: 0.
2:0	DIC: driver impedance control. BIOS: See 2.10.5.4.

F[10:3]RC10 DDR3 Buffer Function [10:3] Control Word 10

F[10:3] corresponds to physical rank [7:0].

Bits	Description
3:2	Reserved.
1:0	QxODT: QxODT. BIOS: See 2.10.5.4.

F[10:3]RC11 DDR3 Buffer Function [10:3] Control Word 11

F[10:3] corresponds to physical rank [7:0].

Bits	Description
3:2	Reserved.
1:0	QxODT: QxODT. BIOS: See 2.10.5.4.

F[11:3]RC12 DDR3 Buffer Function [11:3] Control Word 12

Bits	Description
3:0	Reserved.

F[11:3]RC13 DDR3 Buffer Function [11:3] Control Word 13

Bits	Description
3:0	Reserved.

F3RC[15:14] DDR3 Buffer Function 3 Control Word [15:14]

Bits	Description
3:0	Reserved.

F[11:4]RC[15:14,9:8,6:0] DDR3 Buffer Function [11:4] Control Word [15:14,9:8,6:0]

Bits	Description
3:0	Reserved.

F11RC[11:10] DDR3 Buffer Function 11 Control Word [11:10]

Bits	Description
3:0	Reserved.

F12RC[15:8,6:0] DDR3 Buffer Function 12 Control Word [15:8,6:0]

Bits	Description
3:0	Reserved.

F13RC[9:8,6:0] DDR3 Buffer Function 13 Control Word [9:8,6:0]

Bits	Description
3:0	Reserved.

F13RC10 DDR3 Buffer Function 13 Control Word 10See [2.10.5.7.1.2.3 \[LRDIMM Extended Control Words\]](#).

Bits	Description
3:0	AddressPort[3:0]: address port[3:0].

F13RC11 DDR3 Buffer Function 13 Control Word 11See [2.10.5.7.1.2.3 \[LRDIMM Extended Control Words\]](#).

Bits	Description
3:0	AddressPort[7:4]: address port[7:4].

F13RC[13:12] DDR3 Buffer Function 13 Control Word [13:12]

Bits	Description
3:0	Reserved.

F13RC14 DDR3 Buffer Function 13 Control Word 14See [2.10.5.7.1.2.3 \[LRDIMM Extended Control Words\]](#).

Bits	Description
3:0	DataPort[3:0]: data port[3:0].

F13RC15 DDR3 Buffer Function 13 Control Word 15See [2.10.5.7.1.2.3 \[LRDIMM Extended Control Words\]](#).

Bits	Description
3:0	DataPort[7:4]: data port[7:4].

F[15:14]RC[15:8,6:0] DDR3 Buffer Function [15:14] Control Word [15:8,6:0]

Bits	Description
3:0	Personality: personality. BIOS: See 2.10.5.4 .

2.10.5.7.1.2.3 LRDIMM Extended Control Words

LRDIMM buffer devices contain an extended control word register space. Each control word is eight bits. These registers are programmed by using a series of RCW commands as follows:

1. Write **F13RC10** = addr[3:0].
2. Write **F13RC11** = addr[7:4].
3. Write **F13RC14** = data[3:0].
4. Write **F13RC15** = data[7:4].

ExtRC_xAC DDR3 Buffer Extended Control Word ACh

Bits	Description
7:6	Reserved.
5:4	MRSBroadcastControl: MRS broadcast control. BIOS: 00b.
3	HostBusWriteLevelingControl: host bus write leveling control. BIOS: 0.
2:1	HostMRSSnoopingForwarding: host MRS snooping forwarding. BIOS: 00b.
0	MRSRegisterSource: MRS register source. BIOS: 0.

ExtRC_x[BF:B8] DDR3 Buffer Extended Control Word BFh:B8h

Table 95: Register Mapping for ExtRC_x[BF:B8]

Register	Function
ExtRC_xB8	Physical Rank 0
ExtRC_xB9	Physical Rank 1
ExtRC_xBA	Physical Rank 2
ExtRC_xBB	Physical Rank 3
ExtRC_xBC	Physical Rank 4
ExtRC_xBD	Physical Rank 5
ExtRC_xBE	Physical Rank 6
ExtRC_xBF	Physical Rank 7

Bits	Description
7:6	RttWR: RttWR. BIOS: MR2_dct[1:0][RttWr].
5	Reserved.
4:2	RttNom: RttNom. BIOS: MR1_dct[1:0][RttNom].
1:0	DIC: output driver impedance control. BIOS: MR1_dct[1:0][DIC].

ExtRC_xC8 DDR3 Buffer Extended Control Word C8h

Bits	Description
7	TM: test mode. BIOS: 0.
6:3	CL: CAS latency. BIOS: MR0_dct[1:0][CL].
2	RBT: read burst type. BIOS: MR0_dct[1:0][RBT].
1:0	BL: burst length. BIOS: MR0_dct[1:0][BL].

ExtRC_xC9 DDR3 Buffer Extended Control Word C9h

Bits	Description
7:5	Reserved.
4	PPD: DLL control for precharge powerdown. BIOS: MR0_dct[1:0][PPD].
3:1	WR: write recovery for autoprecharge. BIOS: MR0_dct[1:0][WR].
0	DLL: DLL reset. BIOS: 0.

ExtRC_xCA DDR3 Buffer Extended Control Word CAh

Bits	Description
7:5	Reserved.
4:3	AL: additive latency. BIOS: MR1_dct[1:0][AL].
2:1	Reserved.
0	DLL: DLL enable. BIOS: MR1_dct[1:0][DLL].

ExtRC_xCB DDR3 Buffer Extended Control Word CBh

Bits	Description
7:4	Reserved.
3	TDQS: TDQS enable. BIOS: MR1_dct[1:0][TDQS].
2:0	Reserved.

ExtRC_xCC DDR3 Buffer Extended Control Word CCh

Bits	Description
7	SRT: self refresh temperature range. BIOS: MR2_dct[1:0][SRT]
6	ASR: auto self refresh. BIOS: MR2_dct[1:0][ASR].
5:3	CWL: CAS write latency. BIOS: MR2_dct[1:0][CWL].
2:0	PASR: partial array self refresh. BIOS: MR2_dct[1:0][PASR].

ExtRC_xCD DDR3 Buffer Extended Control Word CDh

Bits	Description
7:0	Reserved.

ExtRC_xCE DDR3 Buffer Extended Control Word CEh

Bits	Description
7:3	Reserved.
2	MPR: MPR operation. BIOS: 0.
1:0	MPRLoc: MPR location. BIOS: 0.

ExtRC_xCF DDR3 Buffer Extended Control Word CFh

Bits	Description
7:0	Reserved.

2.10.5.8 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 [2.10.5.7 \[DRAM Device and Controller Initialization\]](#).

Some of the DRAM training steps described in this section require two passes if the target MEMCLK frequency is greater than the lowest supported MEMCLK frequency. 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\[1:0\]\[MemClkFreq\]](#).

See [2.10.5.6 \[DCT Training Specific Configuration\]](#) 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.

2.10.5.8.1 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 DIMM basis. If the target frequency is greater than the lowest supported MEMCLK frequency then BIOS performs multiple passes; otherwise, only one pass is required. See [2.10.5.3.2.1 \[Requirements for DRAM Frequency Change During Training\]](#).

- Pass 1: Configure the memory subsystem for the lowest supported MEMCLK frequency.
- Pass 2 - Pass N: Configure the memory subsystem for the next higher supported MEMCLK frequency. Repeat until the target MEMCLK frequency is reached.

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

For each DIMM:

1. For UDIMMs or RDIMMs, prepare the DIMMs for write levelization using DDR3-defined MR commands. For LRDIMMs, prepare the DIMMs via RCW commands. See [2.10.5.7.1.1 \[DDR3 MR Initialization\]](#) and [2.10.5.7.1.2 \[Software Control Word Initialization\]](#).

- In the following steps, a quad rank RDIMM is treated as two DIMMs. The target rank and next subsequent rank of a quad rank DIMM are referred to as the target DIMM. The remaining two ranks are treated as a non-target DIMM. A quad logical rank LRDIMM is trained once, but both sets of delay registers must be programmed.
- A. Prepare the target DIMM for write leveling:
 - For UDIMMs and RDIMMs, configure the output driver and on-die termination of the target DIMM as follows:
 - For the first rank of the target DIMM, program `MR1_dct[1:0][Level]=1` and `MR1_dct[1:0][Qoff] = 0`.
 - For all other ranks of the target DIMM, program `MR1_dct[1:0][Level]=1` and `MR1_dct[1:0][Qoff] = 1`.
 - For two or more DIMMs per channel or one of three DIMMs per channel, program `RttNom` of the target rank to the corresponding specified `RttWr` termination. Otherwise, configure `RttNom` of the target DIMM as normal. See 2.10.5.5.5 [DRAM ODT Control].
 - For LRDIMMs, configure write leveling mode and the on-die termination of the target DIMM as follows:
 - Program `F0RC12 = 1h`.
 - For two or more LRDIMMs per channel, program the buffer `RttNom` to the corresponding specified `RttWr` termination. Otherwise, configure the buffer `RttNom` as normal. See 2.10.5.5.5 [DRAM ODT Control].
 - B. Configure `RttNom` on the non-target DIMMs as normal. See 2.10.5.5.5 [DRAM ODT Control].
2. Wait 40 MEMCLKs.
 3. Configure the phy for write levelization training:
 - A. Program `D18F2x9C_x0000_0008_dct[1:0][WrtLvTrEn]=0`.
 - B. Program `D18F2x9C_x0000_0008_dct[1:0][TrDimmSel]` to specify the target DIMM to be trained.
 - Program `D18F2x9C_x0000_0008_dct[1:0][TrNibbleSel]=0`.
 - For x4 DIMMs, BIOS trains both nibbles of a byte lane by programming `D18F2x9C_x0000_0008_dct[1:0][TrNibbleSel]` to specify the nibble. BIOS repeats steps 3 through 5 and uses the average of the trained values for the delay setting.
 - C. Program `D18F2x9C_x0000_0008_dct[1:0][WrLvOdt]` to the proper ODT settings for the current memory subsystem configuration. See 2.10.5.5.5 [DRAM ODT Control].
 - D. Program `D18F2x9C_x0000_0008_dct[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[52:50]_dct[1:0]` to set the gross and fine delay. See 2.10.5.8.1.1 [Write Leveling Seed Value].
 4. Perform write leveling of the devices on the DIMM:
 - A. Program `D18F2x9C_x0000_0008_dct[1:0][WrtLvTrEn]=1`.
 - B. MFENCE.
 - C. Wait 200 MEMCLKs.
 - D. Program `D18F2x9C_x0000_0008_dct[1:0][WrtLvTrEn]=0`.
 - E. Read from registers `D18F2x9C_x0000_00[52:50]_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][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:
 - If LRDIMM then
 - Program `F0RC12 = 0`.
 - Else

- Configure all ranks of the target DIMM for normal operation setting `MR1_dct[1:0][Qoff, Level]= {0b, 0b}`.
- For a two or more DIMM system, program the `RttNom` value for the target DIMM to the normal operating termination.

For each DIMM:

- BIOS calculates and programs the final saved gross and fine delay values for each lane into `D18F2x9C_x0000_00[4A:30]_dct[1:0]` [DRAM DQS Write Timing].
 - `WrDqsFineDly = PhRecFineDly`.
 - `GrossDly = SeedGross + PhRecGrossDly - SeedPreGross`.
 - The Critical Gross Delay (CGD) is the minimum `GrossDly` of all byte lanes and all DIMMs.
 - If (`CGD < 0`) Then
 - `D18F2xA8_dct[1:0][WrDqDqsEarly] = ABS(CGD)`
 - `WrDqsGrossDly = GrossDly + WrDqDqsEarly`
 - Else
 - `D18F2xA8_dct[1:0][WrDqDqsEarly] = 0`.
 - `WrDqsGrossDly = GrossDly`.

2.10.5.8.1.1 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 +/- 1.20 ns, including PVT and jitter, of the measured clock delay.

1. Calculate the total seed based on the following:
 - Pass 1:
 - `SeedTotal = The seed value found in Table 96-Table 98 + (RDIMM & RC2[AddrCmdPrelaunch] ? 0x10 : 0)`.
 - Pass 2 - Pass N:
 - If (RDIMM) then `RegisterDelay = (RC2[AddrCmdPrelaunch] ? 0x30 : 0x20)`.
Else `RegisterDelay = 0`.
 - `SeedTotalPreScaling = (the total delay values D18F2x9C_x0000_00[4A:30]_dct[1:0] [DRAM DQS Write Timing] from the previous pass of write levelization training) - RegisterDelay - (0x20 * D18F2xA8_dct[1:0][WrDqDqsEarly])`.
 - `SeedTotal = RegisterDelay + FLOOR(SeedTotalPreScaling*(target frequency)/(frequency from previous pass))`.
2. If (`SeedTotal >= 0`) then
`SeedGross = SeedTotal DIV 32`.
`SeedFine = SeedTotal MOD 32`.
 else
`SeedGross = (SeedTotal DIV 32) - 1`.
`SeedFine = (SeedTotal MOD 32) + 32`.
 If (`SeedGross` is odd)
 then `SeedPreGross = 1`
 else `SeedPreGross = 2`.
3. Program `D18F2x9C_x0000_00[52:50]_dct[1:0][PhRecFineDly] = SeedFine`.
4. Program `D18F2x9C_x0000_00[52:50]_dct[1:0][PhRecGrossDly] = SeedPreGross`.

Table 96: BIOS Recommendations for Write Leveling Training Seed (G34r1)

Condition				Seed
DIMM	NumDimmSlots	DdrRate	Channel	
UDIMM	1, 2, 3	667	A, B, C, D	Fh
RDIMM				41h
LRDIMM				0h

Table 97: BIOS Recommendations for Write Leveling Training Seed (C32r1)

Condition				Seed
DIMM	NumDimmSlots	DdrRate	Channel	
SODIMM	1, 2	667	A, B	12h
UDIMM	1, 2, 3	667	A, B	12h
RDIMM				3Eh
LRDIMM				0h

Table 98: BIOS Recommendations for Write Leveling Training Seed (AM3r2)

Condition				Seed
DIMM	NumDimmSlots	DdrRate	Channel	
UDIMM, SODIMM	1, 2	667	A, B	Fh

2.10.5.8.2 DQS Receiver Enable Training

Receiver enable delay training is used to dynamically determine the optimal delay value for [D18F2x9C_x0000_00\[2A:10\]_dct\[1:0\] \[DRAM DQS Receiver Enable Timing\]](#). 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 DIMM, per rank basis. If the target frequency is greater than the lowest frequency supported by both the controller and the DIMMs then BIOS performs two passes; otherwise only one pass is required. See [2.10.5.3.2 \[DRAM Channel Frequency Change\]](#).

- Pass 1: Configure the memory subsystem for the lowest supported 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][TrDimmSel]` to specify the target DIMM to be trained.
 - Program `D18F2x9C_x0000_0008_dct[1:0][TrNibbleSel]=0`.
 - For x4 DIMMs, BIOS trains both nibbles of a byte lane by programming `D18F2x9C_x0000_0008_dct[1:0][TrNibbleSel]` to specify the nibble. BIOS repeats steps 2 through 7 and uses the average of the trained values for the delay setting.
2. For each lane program an initial value to registers `D18F2x9C_x0000_00[52:50]_dct[1:0]` to set the gross and fine delay as specified in 2.10.5.8.2.1 [DQS Receiver Enable Training Seed Value].
3. Program `D18F2x9C_x0000_0008_dct[1:0][DqsRcvTrEn]=1`.
4. Issue 192 read requests to the target rank. See 2.10.5.8.6 [Continuous Pattern Generation]. To achieve this, BIOS programs the following:
 - `D18F2x260_dct[1:0][CmdCount] = 192`
 - `D18F2x250_dct[1:0][CmdTgt] = 00b`
 - `D18F2x25[8,4]_dct[1:0][TgtChipSelect] = target rank`
 - `D18F2x25[8,4]_dct[1:0][TgtBank] = 0`
 - `D18F2x25[8,4]_dct[1:0][TgtAddress] = 0`
5. Program `D18F2x9C_x0000_0008_dct[1:0][DqsRcvTrEn]=0`.
6. Read `D18F2x9C_x0000_00[52:50]_dct[1:0][PhRecGrossDly, PhRecFineDly]` 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][DqsRcvEnGrossDelay, DqsRcvEnFineDelay]`. Save the result for use later.
 - $DqsRcvEnFineDelay = PhRecFineDly$.
 - $DqsRcvEnGrossDelay = SeedGross + PhRecGrossDly - SeedPreGross + 1$.
- For each rank pair on a dual-rank or quad-rank DIMM, compute the average value of the total delays saved during the training of each rank and program the result in `D18F2x9C_x0000_00[2A:10]_dct[1:0][DqsRcvEnGrossDelay, DqsRcvEnFineDelay]`.

2.10.5.8.2.1 DQS Receiver Enable Training Seed Value

The seed value for pass 1 of receiver enable delay training is design and platform specific. The seed value represents the total delay from a reference point to the left edge of the read preamble 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. 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:
 - $SeedTotal = \text{The seed value found in Table 99-Table 101} + \text{the total delay value obtained from the first pass of write levelization training. See 2.10.5.8.1 [Write Levelization Training].}$
 - Pass 2 - Pass N:
 - If (RDIMM) then $RegisterDelay = (RC2[AddrCmdPrelaunch] ? 0x30 : 0x20)$.
Else If (LRDIMM) then $RegisterDelay = (F0RC2[AddrCmdPrelaunch] ? 0x30 - (2 * F1RC12[QCA-PrelaunchDelay] : 0x20) + 0x10)$.
Else $RegisterDelay = 0$.
 - $SeedTotalPreScaling = (\text{the total delay values in } D18F2x9C_x0000_00[2A:10]_dct[1:0] \text{ from pass 1 of DQS receiver enable training}) - RegisterDelay - 20h$.
 - $SeedTotal = RegisterDelay + FLOOR(SeedTotalPreScaling * (\text{target frequency}) / (\text{frequency from previous pass}))$.

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[52:50]_dct[1:0][PhRecFineDly] = SeedFine.
6. Program D18F2x9C_x0000_00[52:50]_dct[1:0][PhRecGrossDly] = SeedPreGross.
7. Program D18F2x9C_x0000_00[2A:10]_dct[1:0][DqsRcvEnGrossDelay] = SeedGross.

Table 99: BIOS Recommendations for Receiver Enable Training Seed (G34r1)

Condition				Seed
DIMM	NumDimmSlots	DdrRate	Channel	
UDIMM	1	667	A	3Eh
			B	38h
			C	37h
			D	31h
	2		A	51h
			B	4Ah
			C	46h
			D	3Fh
	3		A	5Eh
			B	52h
			C	48h
			D	3Ch
RDIMM	1	A	43h	
		B	3Fh	
		C	3Ah	
		D	35h	
	2	A	54h	
		B	4Dh	
		C	45h	
		D	40h	
	3	A	6Bh	
		B	5Eh	
		C	4Bh	
		D	3Dh	
LRDIMM	1	A	132h	
		B	122h	
		C	112h	
		D	102h	

Table 100: BIOS Recommendations for Receiver Enable Training Seed (C32r1)

Condition				Seed
DIMM	NumDimmSlots	DdrRate	Channel	
SODIMM	1, 2	667	A	39h
			B	32h
UDIMM	1, 2	667	A	39h
			B	32h
	3		A	45h
			B	37h
RDIMM	1, 2	667	A	3Fh
			B	3Eh
	3		A	47h
			B	38h
LRDIMM	1	667	A	132h
	2		B	122h

Table 101: BIOS Recommendations for Receiver Enable Training Seed (AM3r2)

Condition				Seed
DIMM	NumDimmSlots	DdrRate	Channel	
SODIMM	1, 2	667	A, B	32h
UDIMM	1, 2	667	A, B	3Ah

2.10.5.8.3 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\[1:0\]](#) [DRAM DQS Receiver Enable Timing] to the middle of the received read preamble using the phy phase results.

For each rank and lane:

1. Program [D18F2x9C_x0D0F_0\[F,8:0\]30_dct\[1:0\]](#)[BlockRxDqsLock] = 1.
2. RxEnOrig = [D18F2x9C_x0000_00\[2A:10\]_dct\[1:0\]](#)[DqsRcvEnGrossDelay, DqsRcvEnFineDelay] result from [2.10.5.8.2 \[DQS Receiver Enable Training\]](#).
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\[1:0\]](#)[DqsRcvEnGrossDelay, DqsRcvEnFineDelay] with the current value.
 - B. Perform [2.10.5.8.4 \[DQS Position Training\]](#).
 - 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\[1:0\]](#)[DqsRcvEnGrossDelay, DqsRcvEnFineDelay] = DqsRcvEnCycle - 10h.
6. Program [D18F2x9C_x0D0F_0\[F,8:0\]30_dct\[1:0\]](#)[BlockRxDqsLock] = 0.

2.10.5.8.4 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\[7:5\]_dct\[1:0\]](#) [DRAM Read DQS Timing] and [D18F2x9C_x0000_0\[3:0\]0\[3:1\]_dct\[1:0\]](#) [DRAM Write Data Timing].

Training is accomplished on a per channel, per rank, and per lane basis. For each rank pair on a dual-rank or quad-rank DIMM, BIOS uses the per-lane mutual passing delay values of each rank to calculate the optimal delay values.

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 [2.10.5.8.6.1 \[DRAM Training Pattern Generation\]](#).

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\[1:0\]](#). See [2.10.5.8.5 \[Calculating MaxRdLatency\]](#).

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

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\[3:1\]_dct\[1:0\]](#) from Wr-DQS to Wr-DQS plus 1 UI, using the Wr-DQS delay value found in [2.10.5.8.1 \[Write Levelization Training\]](#):
 - A. Program the write data delay value for the current lane.
 - B. Write the DRAM training pattern to the test address.
 - C. For each read DQS delay value in [D18F2x9C_x0000_0\[3:0\]0\[7:5\]_dct\[1:0\]](#) from 0 to 1 UI:
 - i. Program the read DQS delay value for the current lane.
 - ii. Read the DRAM training pattern from the test address.
 - iii. Record the result for the current settings as a pass or fail depending if the pattern is read correctly.
3. 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\[7:5\]_dct\[1:0\]](#) with the average value of the smallest and largest delay values in the string of consecutive passing results.
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:
 - If LRDIMM then
 - Program [D18F2x9C_x0000_0\[3:0\]0\[3:1\]_dct\[1:0\]](#) with the smallest delay value in the string of consecutive passing results plus one third of the difference of the largest and the smallest delay values in the string of consecutive passing results.
 - Else
 - Program [D18F2x9C_x0000_0\[3:0\]0\[3:1\]_dct\[1:0\]](#) with the average value of the smallest and largest delay values in the string of consecutive passing results.

See [Figure 10](#).

4. $P = P + 5$
5. If $(D18F2xA8_dct[1:0][SubMemclkRegDly] = 0 \ \& \ D18F2x90_dct[1:0][UnbuffDimm] = 0)$
then $P = P + 2$
6. $P = P + (2 * (D18F2x200_dct[1:0][Tcl] - 1 \text{ clocks}))$
7. $P = P + CEIL(MAX(D18F2x9C_x0000_00[2A:10]_dct[1:0][DqsRcvEnGrossDelay, DqsRcvEnFineDelay] + D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0][RdDqsTime] \text{ PCLKs}))$
 -
8. $P = P + 5$
9. $T = T + 800 \text{ ps}$
10. $N = (P / (\text{MemClkFreq} * 2) + T) * \text{NclkFreq}$
 - See $D18F5x1[6C:60][NbDid, NbFid]$ and $D18F2x88_dct[1:0][\text{MemClkFreq}]$.
11. $N = N - 1$.
12. $D18F2x210_dct[1:0]_nbp[3:0][\text{MaxRdLatency}] = CEIL(N) - 1$

2.10.5.8.5.1 MaxRdLatency Training

After DRAM DQS receiver enable training, BIOS optimizes $D18F2x210_dct[1:0]_nbp[3:0][\text{MaxRdLatency}]$ using the following algorithm. For MaxRdLatency training, BIOS generates a training pattern using continuous read or write data streams. See 2.10.5.8.6.1 [DRAM Training Pattern Generation].

For each channel:

1. Calculate a starting MaxRdLatency delay value by executing the steps in 2.10.5.8.5, excluding steps 4, 8, and 11.
2. Select 32 64-byte aligned test addresses associated with the DIMM that has the worst case $(D18F2x9C_x0000_00[2A:10]_dct[1:0][DqsRcvEnGrossDelay, DqsRcvEnFineDelay] + D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0][RdDqsTime])$ 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 $D18F2x210_dct[1:0]_nbp[3:0][\text{MaxRdLatency}]$ 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. Write to $D18F2x9C_x0000_0050$.
5. Program $D18F2x210_dct[1:0]_nbp[3:0][\text{MaxRdLatency}] = CEIL(\text{current value} + 1 \text{ NCLK} + 1.5 \text{ MEM-CLK})$.

2.10.5.8.6 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.

2.10.5.8.6.1 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 $D18F2x250_dct[1:0][\text{CmdTgt}]$. 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 [2.10.5.6 \[DCT Training Specific Configuration\]](#).
- Ensure DIMMs are configured for 8-beat burst length mode. See [2.10.5.7.1.1 \[DDR3 MR Initialization\]](#).
- Wait for `D18F2x250_dct[1:0][CmdSendInProg] = 0`.
- Program `D18F2x250_dct[1:0][CmdTestEnable] = 1`.
- Send activate commands as appropriate. See [2.10.5.8.6.1.1 \[Activate and Precharge Command Generation\]](#).
- Send read or write commands as desired. See [2.10.5.8.6.1.2 \[Read and Write Command Generation\]](#).
- Send precharge commands as appropriate. See [2.10.5.8.6.1.1 \[Activate and Precharge Command Generation\]](#).
- Program `D18F2x250_dct[1:0][CmdTestEnable] = 0`.

2.10.5.8.6.1.1 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 LRDIMM 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 LRDIMM.

For LRDIMMs, 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]]`.

See [2.10.11.1 \[LRDIMM Rank Multiplication\]](#).

2.10.5.8.6.1.2 Read and Write Command Generation

BIOS performs the following steps for read pattern generation:

- Program `D18F2x27C_dct[1:0]`, `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 2.10.5.8.6.1.3 [PRBS Data Comparison].
 - `CmdTgt` corresponding to `D18F2x25[8,4]_dct[1:0]`.
 - `CmdType` = 000b.
 - `SendCmd` = 1.
- 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.
- 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.

2.10.5.8.6.1.3 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]`, `D18F2x298_dct[1:0]`, and `D18F2x29C_dct[1:0]` and can be masked on per-bit basis by programming `D18F2x274_dct[1:0]`, `D18F2x278_dct[1:0]`, and `D18F2x27C_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 102.

Table 102. 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.

2.10.5.9 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] = 11b.
2. Wait 24 MEMCLKs.
3. Program D18F2x94_dct[1:0][DisDramInterface] = 1.

2.10.5.10 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,8:0]13_dct[1:0][DIIIDisEarlyU] = 1.
5. Program D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0][DIIIDisEarlyL] = 1.
6. D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0][RxDqsUDIIPowerDown] = IF (D18F2x90_dct[1:0][X4Dimm]==0) THEN 1 ELSE 0 ENDIF.
7. D18F2x9C_x0D0F_812F_dct[1:0][PARTri] = ~D18F2x90_dct[1:0][ParEn].
8. D18F2x9C_x0D0F_812F_dct[1:0][Add17Tri, Add16Tri] = {1b, 1b} IF (D18F2x94_dct[1:0][MemClk-Freq] <= 800 MHz && ~(mixed channel of x4 and x8 DIMMs)) THEN Program D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0][EnRxPadStandby] = 1. ELSE Program D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0][EnRxPadStandby] = 0. ENDIF.
9. IF (~(mixed channel of x4 and x8 DIMMs)) THEN Program D18F2x9C_x0000_000D_dct[1:0] as follows:
TxMaxDurDIINoLock = RxMaxDurDIINoLock = 7h.
TxCPUUpdPeriod = RxCPUUpdPeriod = 011b.
TxDLLWakeupTime = RxDLLWakeupTime = 11b.

ELSE

Program D18F2x9C_x0000_000D_dct[1:0][TxMaxDurDII_NoLock, RxMaxDurDII_NoLock, TxCPUUpdPeriod, RxCPUUpdPeriod, TxDLLWakeupTime, RxDLLWakeupTime] = {0, 0, 0, 0, 0, 0}.

ENDIF.

10. Program D18F2x9C_x0D0F_0[F,8:0]30_dct[1:0][PwrDn] to disable unused ECC byte lane.

11. Program D18F2x9C_x0D0F_0[F,8:0]04_dct[1:0][TriDM] as follows:

- IF (LRDIMM & (D18F2x90_dct[1:0][X4Dimm] == 0)) THEN 1
- ELSE IF (DataMaskMbType == 0) THEN 1
- ELSE 0
- ENDIF.

2.10.5.10.1 Platform Specific DRAM Phy Tristate Recommendations

BIOS programs the appropriate tri-state or disable register with the values provided below based on the platform and DIMM population. These tables document settings for motherboards which meet the relevant motherboard design guidelines.

Table 103: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (G34r1,UDIMM)

Condition				D18F2x9C_x0000_000C_dct[1:0]			D18F2x88_dct[1:0]
NumDimmSlots	DIMM0	DIMM1	DIMM2	CKETri	ODTri	ChipSelTri	MemClkDis
	1	SR	-	-	2h	Eh	FEh
1	DR	-	-	0h	Ah	FCh	3Ah
2	NP	SR	-	2h	Dh	FBh	3Dh
2	NP	DR	-	0h	5h	F3h	35h
2	SR	SR	-	2h	Ch	FAh	3Ch
2	SR	DR	-	0h	4h	F2h	34h
2	DR	SR	-	0h	8h	F8h	38h
2	DR	DR	-	0h	0h	F0h	30h
3	NP	NP	SR	2h	Bh	EFh	37h
3	NP	NP	DR	0h	Bh	CFh	27h
3	SR	NP	SR	2h	Ah	EEh	36h
3	SR	NP	DR	0h	Ah	CEh	26h
3	DR	NP	SR	0h	Ah	ECh	32h
3	DR	NP	DR	0h	Ah	CCh	22h

Table 104: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (G34r1, RDIMM)

Condition				D18F2x9C_x0000_000C_dct[1:0]		D18F2x88_dct[1:0]	
NumDimmSlots	DIMM0	DIMM1	DIMM2	CKETri	ODTri	ChipSelTri	MemClkDis
	1	SR	-	-	2h	Eh	FCh
1	DR	-	-	0h	Ah	FCh	3Eh
1	QR	-	-	0h	Ah	CCh	3Eh
2	NP	SR	-	2h	Dh	F3h	3Dh
2	NP	DR	-	0h	5h	F3h	3Dh
2	NP	QR	-	0h	5h	33h	3Dh
2	SR	SR	-	2h	Ch	F0h	3Ch
2	SR	DR	-	0h	4h	F0h	3Ch
2	SR	QR	-	0h	4h	30h	3Ch
2	DR	SR	-	0h	8h	F0h	3Ch
2	DR	DR	-	0h	0h	F0h	3Ch
2	DR	QR	-	0h	0h	30h	3Ch
2	QR	SR	-	0h	8h	C0h	3Ch
2	QR	DR	-	0h	0h	C0h	3Ch
2	QR	QR	-	0h	0h	0h	3Ch
3	NP	NP	SR	2h	Bh	CFh	37h
3	NP	NP	DR	0h	Bh	CFh	37h
3	NP	QR	NP	0h	5h	33h	3Dh
3	NP	QR	SR, DR	0h	1h	3h	35h
3	SR	NP	SR	2h	Ah	CCh	36h
3	SR	NP	DR	0h	Ah	CCh	36h
3	SR	SR	SR	2h	8h	C0h	34h
3	SR	SR	DR	0h	8h	C0h	34h
3	SR, DR	DR	SR, DR	0h	0h	C0h	34h
3	SR, DR	QR	SR, DR	0h	0h	0h	34h
3	DR	NP	SR, DR	0h	Ah	CCh	36h
3	DR	SR	SR, DR	0h	8h	C0h	34h

Table 105: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (G34r1, LRDIMM)

Condition						D18F2x9C_x0000_000C_dct[1:0]	D18F2x88_dct[1:0]		
NumDimmSlots	DIMM0	DIMM1	DIMM2	D18F2xA8_dct[1:0][CsMux67]	D18F2xA8_dct[1:0][CsMux45]	CKETri	ODTTri	ChipSelTri	MemClkDis
	1	SR	-	-	-	-	2h	Eh	FCh
1	DR	-	-	-	0	0h	Eh	FCh	3Eh
1	DR	-	-	-	1	0h	Eh	CCh	3Eh
1	QR	-	-	-	-	0h	Eh	CCh	3Eh
2	NP	SR	-	-	-	2h	Dh	F3h	3Dh
2	NP	DR	-	0	-	0h	Dh	F3h	3Dh
2	NP	DR	-	1	-	0h	Dh	33h	3Dh
2	NP	QR	-	-	-	0h	Dh	33h	3Dh
2	SR	SR	-	-	-	2h	Ch	F0h	3Ch
2	SR	DR	-	0	-	0h	Ch	F0h	3Ch
2	SR	DR	-	1	-	0h	Ch	30h	3Ch
2	SR, DR	QR	-	-	0	0h	Ch	30h	3Ch
2	SR, DR	QR	-	-	1	0h	Ch	0h	3Ch
2	DR	SR, DR	-	0	0	0h	Ch	F0h	3Ch
2	DR	SR, DR	-	0	1	0h	Ch	C0h	3Ch
2	DR	SR, DR	-	1	0	0h	Ch	30h	3Ch
2	DR	SR, DR	-	1	1	0h	Ch	0h	3Ch
2	QR	SR, DR	-	0	-	0h	Ch	C0h	3Ch
2	QR	SR, DR	-	1	-	0h	Ch	00h	3Ch
2	QR	QR	-	-	-	0h	Ch	0h	3Ch
3	NP	NP	SR	-	-	2h	Bh	CFh	37h
3	NP	NP	DR	0	-	0h	Bh	CFh	37h
3	NP	NP	DR	1	-	0h	Bh	0Fh	37h
3	SR	NP	SR	-	-	2h	Ah	CCh	36h
3	SR	NP	DR	0	-	0h	Ah	CCh	36h
3	SR	NP	DR	1	-	0h	Ah	Ch	36h
3	SR	SR	SR	-	-	2h	8h	C0h	34h
3	SR	SR	DR	0	-	0h	8h	C0h	34h
3	SR	SR	DR	1	-	0h	8h	0h	34h
3	SR	DR	SR, DR	0	-	0h	8h	C0h	34h

Table 105: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (G34r1, LRDIMM)

Condition						D18F2x9C_x0000_000C_dct[1:0]		D18F2x88_dct[1:0]	
NumDimmSlots	DIMM0	DIMM1	DIMM2	D18F2xA8_dct[1:0][CsmMux67]		CKETri	ODTTri	ChipSelTri	MemClkDis
3	SR	DR	SR, DR	1	-	0h	8h	0h	34h
3	DR	NP	SR, DR	0	-	0h	Ah	CCh	36h
3	DR	NP	SR, DR	1	-	0h	Ah	Ch	36h
3	DR	SR, DR	SR, DR	0	-	0h	8h	C0h	34h
3	DR	SR, DR	SR, DR	1	-	0h	8h	0h	34h

Table 106: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (C32r1, UDIMM or SODIMM)

Condition				D18F2x9C_x0000_000C_dct[1:0]		D18F2x88_dct[1:0]	
NumDimmSlots	DIMM0	DIMM1	DIMM2	CKETri	ODTTri	ChipSelTri	MemClkDis
1	DR	-	-	0h	Ah	FCh	3Ah
2	NP	SR	-	2h	Dh	FBh	3Dh
2	NP	DR	-	0h	5h	F3h	35h
2	SR	SR	-	2h	Ch	FAh	3Ch
2	SR	DR	-	0h	4h	F2h	34h
2	DR	SR	-	0h	8h	F8h	38h
2	DR	DR	-	0h	0h	F0h	30h
3	NP	NP	SR	2h	Bh	EFh	3Bh
3	NP	NP	DR	0h	Bh	CFh	33h
3	SR	NP	SR	2h	Ah	EEh	3Ah
3	SR	NP	DR	0h	Ah	CEh	32h
3	DR	NP	SR	0h	Ah	ECh	38h
3	DR	NP	DR	0h	Ah	CCh	30h

Table 107: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (C32r1, RDIMM)

Condition				D18F2x9C_x0000_000C_dct[1:0]	D18F2x88_dct[1:0]		
NumDimmSlots	DIMM0	DIMM1	DIMM2	CKETri	ODTri	ChipSelTri	MemClkDis
	1	SR	-	-	2h	Eh	FCh
1	DR	-	-	0h	Ah	FCh	3Eh
1	QR	-	-	0h	Ah	CCh	3Eh
2	NP	SR	-	2h	Dh	F3h	3Dh
2	NP	DR	-	0h	5h	F3h	3Dh
2	NP	QR	-	0h	5h	33h	3Dh
2	SR	SR	-	2h	Ch	F0h	3Ch
2	SR	DR	-	0h	4h	F0h	3Ch
2	SR	QR	-	0h	4h	30h	3Ch
2	DR	SR	-	0h	8h	F0h	3Ch
2	DR	DR	-	0h	0h	F0h	3Ch
2	DR	QR	-	0h	0h	30h	3Ch
2	QR	SR	-	0h	8h	C0h	3Ch
2	QR	DR	-	0h	0h	C0h	3Ch
2	QR	QR	-	0h	0h	0h	3Ch
3	NP	NP	SR	2h	Bh	CFh	3Bh
3	NP	NP	DR	0h	Bh	CFh	3Bh
3	NP	QR	NP	0h	5h	33h	3Dh
3	NP	QR	SR, DR	0h	1h	3h	39h
3	SR	NP	SR	2h	Ah	CCh	3Ah
3	SR	NP	DR	0h	Ah	CCh	3Ah
3	SR	SR	SR	2h	8h	C0h	38h
3	SR	SR	DR	0h	8h	C0h	38h
3	SR, DR	DR	SR, DR	0h	0h	C0h	38h
3	SR, DR	QR	SR, DR	0h	0h	0h	38h
3	DR	NP	SR, DR	0h	Ah	CCh	3Ah
3	DR	SR	SR, DR	0h	8h	C0h	38h

Table 108: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (C32r1, LRDIMM)

Condition						D18F2x9C_x0000_000C_dct[1:0]	D18F2x88_dct[1:0]		
NumDimmSlots	DIMM0	DIMM1	DIMM2	D18F2xA8_dct[1:0][CsMux67]	D18F2xA8_dct[1:0][CsMux45]	CKETri	ODTTri	ChipSelTri	MemClkDis
	1	SR	-	-	-	-	2h	Eh	FCh
1	DR	-	-	-	0	0h	Eh	FCh	3Eh
1	DR	-	-	-	1	0h	Eh	CCh	3Eh
1	QR	-	-	-	-	0h	Eh	CCh	3Eh
2	NP	SR	-	-	-	2h	Dh	F3h	3Dh
2	NP	DR	-	0	-	0h	Dh	F3h	3Dh
2	NP	DR	-	1	-	0h	Dh	33h	3Dh
2	NP	QR	-	-	-	0h	Dh	33h	3Dh
2	SR	SR	-	-	-	2h	Ch	F0h	3Ch
2	SR	DR	-	0	-	0h	Ch	F0h	3Ch
2	SR	DR	-	1	-	0h	Ch	30h	3Ch
2	SR, DR	QR	-	-	0	0h	Ch	30h	3Ch
2	SR, DR	QR	-	-	1	0h	Ch	0h	3Ch
2	DR	SR, DR	-	0	0	0h	Ch	F0h	3Ch
2	DR	SR, DR		0	1	0h	Ch	C0h	3Ch
2	DR	SR, DR		1	0	0h	Ch	30h	3Ch
2	DR	SR, DR		1	1	0h	Ch	0h	3Ch
2	QR	SR, DR	-	0	-	0h	Ch	C0h	3Ch
2	QR	SR, DR	-	1	-	0h	Ch	0h	3Ch
2	QR	QR	-	-	-	0h	Ch	0h	3Ch
3	NP	NP	SR	-	-	2h	Bh	CFh	3Bh
3	NP	NP	DR	0	-	0h	Bh	CFh	3Bh
3	NP	NP	DR	1	-	0h	Bh	Fh	3Bh
3	SR	NP	SR	-	-	2h	Ah	CCh	3Ah
3	SR	NP	DR	0	-	0h	Ah	CCh	3Ah
3	SR	NP	DR	1	-	0h	Ah	Ch	3Ah
3	SR	SR	SR	-	-	2h	8h	C0h	38h
3	SR	SR	DR	0	-	0h	8h	C0h	38h
3	SR	SR	DR	1	-	0h	8h	0h	38h
3	SR	DR	SR, DR	0	-	0h	8h	C0h	38h

Table 108: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (C32r1, LRDIMM)

Condition						D18F2x9C_x0000_000C_dct[1:0]		D18F2x88_dct[1:0]	
NumDimmSlots	DIMM0	DIMM1	DIMM2	D18F2xA8_dct[1:0][CsMux67]	D18F2xA8_dct[1:0][CsMux45]	CKETri	ODTTri	ChipSelTri	MemClkDis
	3	SR	DR	SR, DR	1	-	0h	8h	0h
3	DR	NP	SR, DR	0	-	0h	Ah	CCh	3Ah
3	DR	NP	SR, DR	1	-	0h	Ah	Ch	3Ah
3	DR	SR, DR	SR, DR	0	-	0h	8h	C0h	38h
3	DR	SR, DR	SR, DR	1	-	0h	8h	0h	38h

Table 109: BIOS Recommendations for CKETri, ODTTri, ChipSelTri, MemClkDis (AM3r2, UDIMM or SODIMM)

Condition			D18F2x9C_x0000_000C_dct[1:0]		D18F2x88_dct[1:0]	
NumDimmSlots	DIMM0	DIMM1	CKETri	ODTTri	ChipSelTri	MemClkDis
	1	SR	-	2h	Eh	FEh
1	DR	-	0h	Ah	FCh	3Ch
2	NP	SR	2h	Dh	FBh	37h
2	NP	DR	0h	5h	F3h	33h
2	SR	SR	2h	Ch	FAh	35h
2	SR	DR	0h	4h	F2h	31h
2	DR	SR	0h	8h	F8h	34h
2	DR	DR	0h	0h	F0h	30h

2.10.6 Memory Interleaving Modes

The processor supports three 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 2.10.6.1 [Chip Select Interleaving].
- Channel: interleaves the physical address space over multiple channels, as opposed to each channel owning single consecutive address spaces. See 2.10.6.2 [Channel Interleaving].
- Node: interleaves the physical address space over multiple nodes, as opposed to each node owning single consecutive address spaces. See 2.10.6.3 [Node Interleaving].

Any combination of these interleaving modes may be enabled concurrently. In addition, it is possible to remap the physical addresses corresponding to an interleaved region of memory. See [Table 110](#) and [D18F2x10C](#).

Table 110. Recommended interleave configurations

Interleaving Mode	Enabled	Disabled
Chip Select Interleaving	Number of chip selects installed on the channel is a power of two.	Requirements not satisfied.
Channel Interleaving	DIMMs are present on both channels.	Requirements not satisfied.
Node Interleaving	Programmed specific to the current platform configuration.	Programmed specific to the current platform configuration or requirements not satisfied.
Interleave Region Remapping ¹	UMA and DIMMs are present on both channels and the channels do not have the same amount of DRAM.	~UMA or both channels have the same amount of DRAM or DIMMs are not present on both channels.
1. The channel interleave region should always include the frame buffer.		

2.10.6.1 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[21:13] bits as defined in [Table 111](#).
3. For each enabled chip select, swap the corresponding AddrMask[38:27] bits with the AddrMask[21:13] bits as defined in [Table 111](#).

Table 111. DDR3 Swapped Normalized Address Lines for CS Interleaving

DIMM Address Map ¹	Chip Select Size ²	Swapped Base Address and Address Mask bits ²		
		8 way CS interleaving	4 way CS interleaving	2 way CS interleaving
0001b	256-MB	[30:28] and [18:16]	[29:28] and [17:16]	[28] and [16]
0010b	512-MB	[31:29] and [18:16]	[30:29] and [17:16]	[29] and [16]
0101b	1-GB, 2-GB, 4-GB	[32:30] and [18:16]	[31:30] and [17:16]	[30] and [16]
0110b	1-GB	[32:30] and [19:17]	[31:30] and [18:17]	[30] and [17]
0111b	2-GB, 4-GB, 8-GB	[33:31] and [18:16]	[32:31] and [17:16]	[31] and [16]
1000b	2-GB, 4-GB, 8-GB	[33:31] and [19:17]	[32:31] and [18:17]	[31] and [17]

Table 111. DDR3 Swapped Normalized Address Lines for CS Interleaving

DIMM Address Map ¹	Chip Select Size ²	Swapped Base Address and Address Mask bits ²		
		8 way CS interleaving	4 way CS interleaving	2 way CS interleaving
1001b	4-GB, 8-GB, 16-GB	[34:32] and [19:17]	[33:32] and [18:17]	[32] and [17]
1010b	4-GB, 8-GB, 16-GB	[34:32] and [18:16]	[33:32] and [17:16]	[32] and [16]
1011b	8-GB, 16-GB, 32-GB	[35:33] and [19:17]	[34:33] and [18:17]	[33] and [17]
1100b	16-GB	[36:34] and [20:18]	[35:34] and [19:18]	[34] and [18]

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

2. For LRDIMMs BIOS adds an offset to the bit positions specified based on [D18F2x\[6C:60\]_dct\[1:0\]\[RankDef\]](#) as follows:
 RankDef=0xb: 0
 RankDef=10b: 1
 RankDef=11b: 2

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_3FE0h. // CS0/CS1 = 256 MB
 - [D18F2x64](#) = 0008_3FE0h. // CS2/CS3 = 256 MB
- The base address bits to be swapped are defined in [Table 111](#), 256MB chip select size, 4 way CS interleaving column. The BaseAddr[29:28] bits are specified by [D18F2x\[5C:40\]_dct\[1:0\]\[BaseAddr\[21:20\]\]](#). The BaseAddr[17:16] bits are specified by [D18F2x\[5C:40\]_dct\[1:0\]\[BaseAddr\[9:8\]\]](#).
 - [D18F2x40](#) = 0000_0001h.
 - [D18F2x44](#) = 0000_0101h.
 - [D18F2x48](#) = 0000_0201h.
 - [D18F2x4C](#) = 0000_0301h.
- 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\[1:0\]\[AddrMask\[21:20\]\]](#). The AddrMask[17:16] bits are specified by [D18F2x\[6C:60\]_dct\[1:0\]\[AddrMask\[9:8\]\]](#).
 - [D18F2x60](#) = 0038_3CE0h.
 - [D18F2x64](#) = 0038_3CE0h.

2.10.6.2 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 [2.10.7 \[Memory Hoisting\]](#).

2.10.6.3 Node Interleaving

The node memory interleaving mode has the following requirements:

- All nodes in the system must contain the same amount of DRAM.
- All the DRAM of all nodes in the system must be node interleaved.
- The base and limit registers for all nodes must be programmed to 0 and top of memory, respectively.
- All DRAM channels in the system must have the same amount of DRAM on the respective channel across all nodes.
- `D18F2x110[DctSelIntLvAddr]==01b` not supported.

Node interleaving for up to 8 nodes is controlled by `D18F1x[17C:140,7C:40][IntlvEn and IntlvSel]`, `D18F1x120[DramIntlvSel]` and `D18F1x124[DramIntlvEn]`. `IntlvEn` and `DramIntlvEn` are programmed to specify the interleaved address bits (programmed the same in each node). `IntlvSel` specifies the value that those address bits need to be to target a node (must be programmed to a different value for each node). `DramIntlvSel` specifies the value of those address bits for the local node. It is expected that one `D18F1x[17C:140,7C:40][DRAM Base/Limit]` pair is enabled per node; one of these pairs selects the local node by having an `IntlvSel` value that matches `D18F1x120[DramIntlvSel]`; `IntlvEn` is the same in all `D18F1x[17C:140,7C:40][DRAM Base/Limit]` pairs and the same as `D18F1x124[DramIntlvEn]` of all nodes. For example, a 4-node system would normally be programmed as follows for interleaving:

Node 0 - <code>IntlvEn = 011b</code> - <code>IntlvSel = 000b</code> - <code>Addr[13:12] = 00b</code>	Node 1 - <code>IntlvEn = 011b</code> - <code>IntlvSel = 001b</code> - <code>Addr[13:12] = 01b</code>
Node 2 - <code>IntlvEn = 011b</code> - <code>IntlvSel = 010b</code> - <code>Addr[13:12] = 10b</code>	Node 3 - <code>IntlvEn = 011b</code> - <code>IntlvSel = 011b</code> - <code>Addr[13:12] = 11b</code>

2.10.7 Memory Hoisting

Memory hoisting reclaims the otherwise inaccessible DRAM that would naturally reside in memory regions used by MMIO and HyperTransport™. 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. 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 `DramHoleSize[31:24] = 100h - D18F1xF0[DRAMHoleBase[31:24]]`.

2.10.7.1 DramHoleOffset Programming

D18F1xF0[DramHoleOffset] is programmed based on the scenarios shown in Figure 12:

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

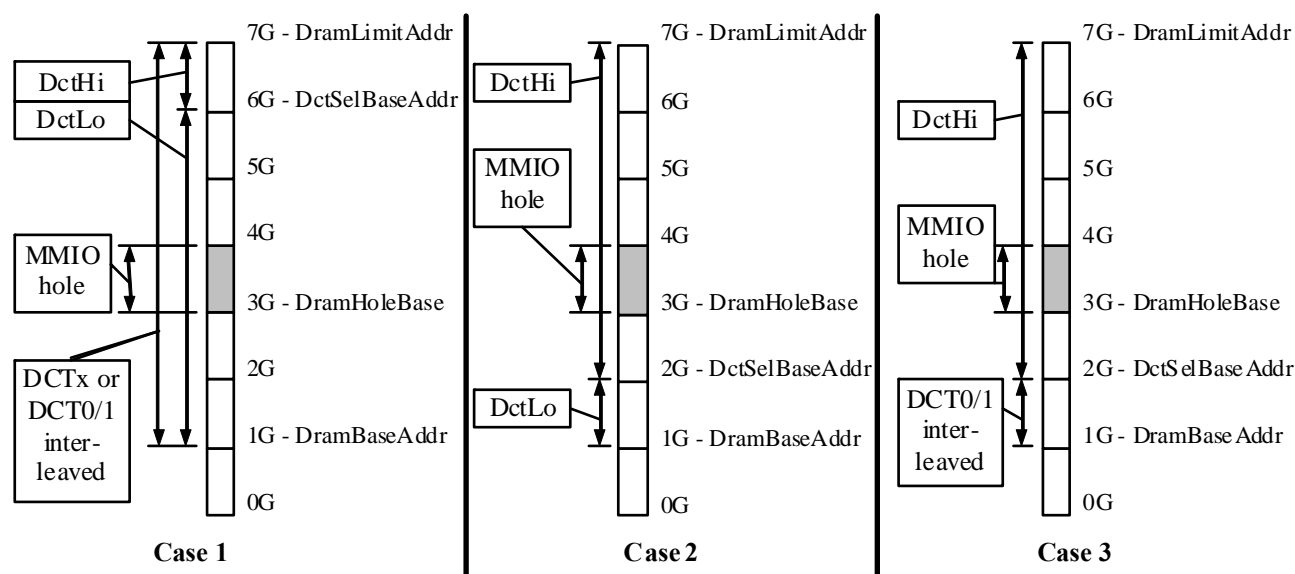


Figure 12: Example Cases for Programming DramHoleOffset

2.10.7.2 DctSelBaseOffset Programming

When D18F2x110[DctSelHiRngEn] = 1, D18F2x114[DctSelBaseOffset] is programmed based on the scenarios shown in Figure 13:

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

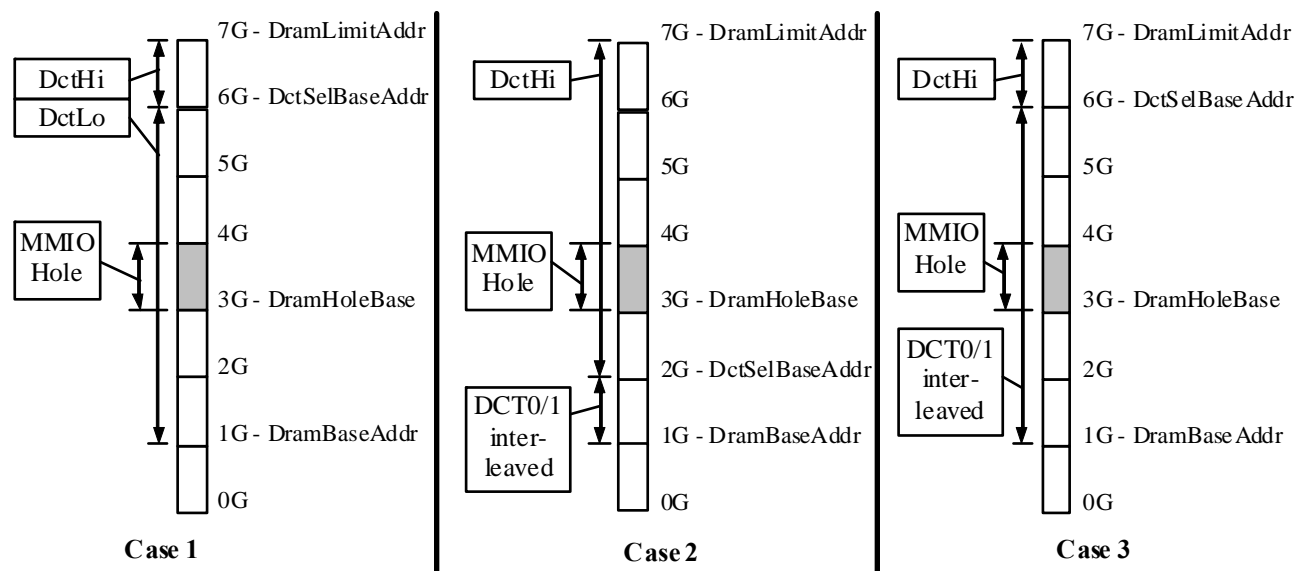


Figure 13: Example Cases for Programming DctSelBaseOffset

2.10.8 DRAM CC6 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 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 on a node is defined to be a fixed 16MB.

For single node systems, BIOS configures the storage region at the top of the DRAM range and adjusts [D18F1x\[17C:140,7C:40\]\[DramLimit\]](#) downward accordingly. BIOS sets [D18F4x128\[CoreStateSaveDestNode\] = D18F0x60\[NodeId\]](#).

For multiple node systems, BIOS sets [D18F4x128\[CoreStateSaveDestNode\]](#) equal to the [D18F0x60\[NodeId\]](#) of the node at the top of the DRAM address map and adjusts [D18F1x\[17C:140,7C:40\]\[DramLimit\]](#) downward accordingly. If node interleaving is enabled by [D18F1x\[17C:140,7C:40\]\[IntlvEn\]](#), the amount of system memory available is reduced by 16MB * the number of nodes. In this case, BIOS sets [D18F4x128\[CoreStateSaveDestNode\] = D18F0x60\[NodeId\]](#) for each node.

See [Table 112](#), [Table 113](#), and [Table 114](#).

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

Table 112. Single Node Example

Node	DRAM Populated	D18F1x[17C:140,7C:40][DramBase, DramLimit]	CC6 DRAM Range	D18F4x128[CoreStateSaveDestNode]	D18F1x120[DramBaseAddr], D18F1x124[DramLimitAddr]
0	256 MB	0 MB, 240 MB - 1	240 MB, 256 MB - 1	0	0 MB, 256 MB - 1

Table 113. Dual Node Example

Node	DRAM Populated	D18F1x[17C:140,7C:40] [DramBase, DramLimit]	CC6 DRAM Range	D18F4x128 [CoreStateSa veDestNode]	D18F1x120[DramBaseAddr], D18F1x124[DramLimitAddr]
0	256 MB	0 MB, 240 MB - 1	240 MB, 256 MB - 1	0	0 MB, 256 MB - 1
1	0 MB	-	-	0	-

Table 114. Dual Node with Node Interleaved Example

Node	DRAM Populated	D18F1x[17C:140,7C:40] [DramBase, DramLimit]	CC6 DRAM Range	D18F4x128 [CoreStateSa veDestNode]	D18F1x120[DramBaseAddr], D18F1x124[DramLimitAddr]
0	256 MB	0 MB, 480 MB - 1	480 MB, 512 MB - 1	0	0 MB, 512 MB - 1
1	256 MB	0 MB, 480 MB - 1	480 MB, 512 MB - 1	1	0 MB, 512 MB - 1

2.10.9 On-Line Spare

On-line spare is a RAS mechanism that allows the system to reserve one rank of a DIMM to be used as a spare rank. System software reserves a spare rank (logical rank for LRDIMMs) by setting [D18F2x\[5C:40\]_dct\[1:0\]\[Spare\]](#) in one of the CS Base address registers. [D18F2x80_dct\[1:0\]](#) and [D18F2x\[6C:60\]_dct\[1:0\]\[RankDef\]](#) are configured specific to the spare DIMM. The spare rank must be greater than or equal to the size of all other ranks on the channel.

The system can switch to the spare rank when system software determines that one of the ranks in the system is no longer functioning properly and needs to be replaced. The on-line spare mechanism is controlled by [D18F3xB0 \[On-Line Spare Control\]](#). System software initiates the swap to the spare rank by writing the chip select number of the bad rank to [D18F3xB0\[BadDramCS1, BadDramCS0\]](#) and setting [D18F3xB0\[SwapEn1, SwapEn0\]](#). The spare rank inherits in hardware the [D18F2x\[5C:40\]_dct\[1:0\]\[BaseAddr\]](#) and [D18F2x\[6C:60\]_dct\[1:0\]\[AddrMask\]](#) values from the bad rank.

On-line spare is not supported in UMA systems.

2.10.9.1 On-Line Spare and CS Interleaving

The on-line spare feature can only be used with 2 way and 4 way CS interleaving, under the following conditions.

- All ranks of each DIMM (logical ranks for LRDIMMs) present must be of the same size and configuration.
- [Table 115](#) lists the supported DIMM populations.
 - BIOS sets Spare on one CS of one of the DIMMs with smallest number of ranks. When the TestFail column is indicated, BIOS sets TestFail on the CS that shares a CS Mask register with the spare rank.

Table 115. On-Line Spare and CS Interleaving Configurations

DIMMs/Ch	SR	DR	QR	CS Interleave Ways	TestFail
1	-	-	1	2	1
2	1	1	-	2	-
	1	-	1	4	-
	-	2	-	2	1
	-	1	1	4	1
3	3	-	-	2	-
	1	2	-	4	-
	-	3	-	4	1

2.10.10 DRAM On DIMM Thermal Management and Power Capping

Each DCT can throttle commands based on the state of the channel EVENT_L pin (G34r1 and C32r1 packages only) 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 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.
 - BIOS programs `D18F2xA4[ODTSEn, CmdThrottleMode]` only once immediately following DRAM training.
 - 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).
 - On a channel basis, 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.
 - On a channel basis, 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

`D18F2xA4[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 guard-band 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.

APML may enable bandwidth capping on a DRAM controller by setting `MSRC001_0079[BwCapEn] = 1` and programming `MSRC001_0079[BwCapCmdThrottleMode]` with the throttling mode to employ. The DCT will employ the larger of the two throttling percentages as specified by `MSRC001_0079[BwCapCmdThrottleMode]` and `D18F2xA4[CmdThrottleMode]` if the EVENT_L pin is asserted when both `D18F2xA4[BwCapEn]==1` and `D18F2xA4[ODTSEn]==1`.

2.10.11 LRDIMM

Processor revisions support JEDEC DDR3 LRDIMMs, which provide load isolation on the DRAM address, command, and data buses, and allows the DCT to support larger DRAM capacities.

2.10.11.1 LRDIMM Rank Multiplication

LRDIMMs support a rank multiplication feature which allows more than 4 physical ranks per-DIMM while using a RDIMM compatible connector. Physical ranks are accessed on the LRDIMM by specifying a logical rank of either S0 or S1 and then selecting a physical rank via extended address bits on MA[15:14], S2, and S3 depending on the DRAM devices.

The DCT supports LRDIMM rank multiplication modes using extended row address bits as shown in Table 164. BIOS enables rank multiplication mode by programming `D18F2x[6C:60]_dct[1:0][RankDef]`. `D18F2xA8_dct[1:0][CsMux67, CsMux45]` must also be configured to agree with the F0RC15 control word and in accordance with the relevant AMD motherboard design guideline.

2.11 Thermal Functions

Thermal functions 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 physical 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.

•

2.11.1 The Tctl Temperature Scale

Tctl is a processor temperature control value used for processor thermal management. Tctl is accessible through `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,max is specified in the power and thermal data sheet. 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 14: Tctl scale

2.11.2 Temperature-Driven Logic

The temperature calculated by the TCC is used by HTC, THERMTRIP, and PROCHOT_L.

2.11.2.1 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 [2.5.2 \[P-states\]](#). 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. The minimum assertion and deassertion time for PROCHOT_L is 200 μ s with a minimum period of 2 ms. See [2.5.2.1.5 \[Core P-state](#)

Limits] and 2.4.8.1 [Local APIC].

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 is deasserted.
- 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 is asserted.

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 is deasserted.

The default value of the HTC temperature threshold (Tctl_max) is specified in the Power and Thermal Data-sheet.

2.11.2.2 Software P-state Limit Control

D18F3x68 [Software P-state Limit] provides a software mechanism to limit the P-state MSRC001_0061 [CurP-stateLimit]. See 2.5.2 [P-states].

2.11.2.3 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.

2.12 Links

A *link* is a block of link signals, including 16 CAD signals, 2 CTL signals, and 2 CLK signals. Links may support ungangned modes in which subgroups of link signals--or *sublinks*--are connected to separate devices, as specified by D18F3xE8 [Northbridge Capabilities][UnGangEn]. Links may operate per coherent protocol or IO protocol. The electrical definition is per various revisions of the link specification; the terminology for these modes is as follows:

- Gen1: refers to link rates of 0.4 to 1.6 GT/s in the revision 1 specification or 2.0 GT/s in the revision 2 specification.
 - **G34r1**: 1.0 GHz disallowed for links connecting internal nodes as indicated by **D18F0x1A0**[IntLnk-Route]. See **D18F0x[E8,C8,A8,88]**[Freq[3:0]].
 - **AM3r2**: 0.4-1.0 GHz disallowed for AM3r2 package. 0.2 GHz supported only as boot frequency to first warm reset; 0.2 GHz disallowed as operational frequency. See **D18F0x[E8,C8,A8,88]**[Freq[3:0]].
- Gen3: refers to link rates of 2.4 to 6.4 GT/s in the revision 3 specification. 2.4 GT/s and 3.2 GT/s are supported as specified in the revision 3 specification only, not as specified in the revision 2 specification.

2.12.1 Link Initialization

2.12.1.1 Ganging And Unganging

The following combinations of maximum bit widths (it is always possible to connect to a device using a supported, narrower bit width), protocols, and frequencies are supported:

- One 16-bit link (ganged); either IO or coherent protocol; any supported link frequency. In ganged mode, the link may or may not be left unconnected. In ganged mode, registers that control sublink 0 control the entire link; registers that control sublink 1 are reserved.
- Two 8-bit links (unganged); the two sublinks may be configured for any combination of IO or coherent protocol; if the two link frequencies are the same, then they may be any supported frequency; if the two link frequencies are different, then they are required to be one of the following ratios to each other: 8:1, 6:1, 4:1, 2:1; legal combinations are {5.6, 2.8}, {6.4, 3.2}, {6.4, 1.6}, {6.4, 0.8}, {2.4, 0.4}, {4.8, 0.8}, {4.8, 2.4}, {4.8, 1.2}, {4.0, 2.0}, {3.2, 1.6}, {3.2, 0.8}, {3.2, 0.4}, {2.4, 1.2}, {1.6, 0.8}, {1.6, 0.4}, and {0.8, 0.4} GT/s). The following combinations are allowed for test and debug purposes: {6.4, 0.8}, {2.4, 0.4}, {4.8, 0.8}, {3.2, 0.8}, {3.2, 0.4}, {1.6, 0.8}, and {1.6, 0.4}. In unganged mode, neither, either, or both of the two 8-bit sublinks may be left unconnected. In unganged mode, sublink 0 refers to the link associated with **CLK[0]**, **CTL[0]**, and **CAD[7:0]**; sublink 1 refers to the link associated with **CLK[1]**, **CTL[1]**, and **CAD[15:8]**.

2.12.1.2 Ganging Detection And Control

If unganging is not supported by the processor (**D18F3xE8**[UnGangEn]), then the links always cold boot to the ganged state.

Otherwise, the ganged state of the links at cold boot is based on the state of **CTL[1]**. If **CTL[1]=0**, then the link powers up unganged. If **CTL[1]=1**, then the link powers up ganged. If **CTL[1]** is connected between the processor and another device (such as another processor) that supports the Gen3 link specification, then the link cold boots to the unganged state.

If both sublinks of an unganged link connect the same two devices, then initialization software may be used to place these sublinks into the ganged state (**D18F0x[18C:170]**[Ganged]).

Unused upper sublinks should be left disconnected with processor inputs floating to minimize power consumption. Software should not gang inactive sublinks with active sublinks.

2.12.1.3 Link Type Detect

The link may be initialized in one of the following states during cold reset:

- The link may be ganged or unganged.
- The link/sublink is not connected with inputs terminated to the proper state to indicate ganged/unganged.
- The link/sublink is not connected with inputs floating (as with a connection to an unpopulated socket).

The processor follows the protocol described in the Gen3 link specification to determine the cold boot state of

D18F0x[18C:170][Ganged] and D18F0x[E4,C4,A4,84][TransOff, EndOfChain].

D18F0x[E4,C4,A4,84][TransOff, EndOfChain] are set when the link is unconnected, as follows:

- Unterminated link: no device is detected on the other side of the link.
- DC-coupled links: link is strapped in the unconnected state per the link specification.

2.12.1.3.1 Link Specific Registers

The following registers are link specific:

- F[4,0]x[9C:80] are associated with link 0.
- F[4,0]x[BC:A0] are associated with link 1.
- F[4,0]x[DC:C0] are associated with link 2.
- F[4,0]x[FC:E0] are associated with link 3.
- The function 0 registers are associated with the whole link if it is ganged or sublink 0 if it is unganged; the function 4 registers are associated with sublink 1 if the link is unganged.

2.12.1.3.2 Unused Links

When both sublinks of a link are unused and disconnected, the clocks to the link are disabled and the following registers become unavailable. Writes to these registers are ignored and reads are undefined.

- D18F3x1[54,50,4C,48].
- D18F0x[F0,D0,B0,90] and D18F4x[F0,D0,B0,90].
- D18F0x[F4,D4,B4,94] and D18F4x[F4,D4,B4,94].
- D18F4x1[98,90,88,80].
- D18F4x1[9C,94,8C,84].

2.12.1.4 Legal Topologies

The link may be connected in these configurations:

- 16-bit Gen3 device connected (CTL[1] connected)
- 16-bit Gen1 device connected (CTL[1] terminated)
- Unganged:
 - Two 8-bit devices connected
 - One 8-bit device connected to either sublink and the other sublink with inputs terminated
 - One 8-bit device connected to either sublink and the other sublink with inputs floating
- Link inputs terminated
- Link inputs floating

2.12.1.5 Link Mapping Between Package and Node

The AM3r2 package maps package link 0 to node link 0.

The C32r1 package maps package link 0 to node link 0, package link 1 to node link 1, and package link 2 to node link 2.

The G34r1 package only supports dual-node processors (D18F3xE8[MultiNodeCpu]=1). The links on the G34r1 package are connected to the links of the internal nodes as shown in [Figure 15](#). See [2.9.3.2 \[Hyper-Transport™ Technology Routing\]](#) for supported system topologies.

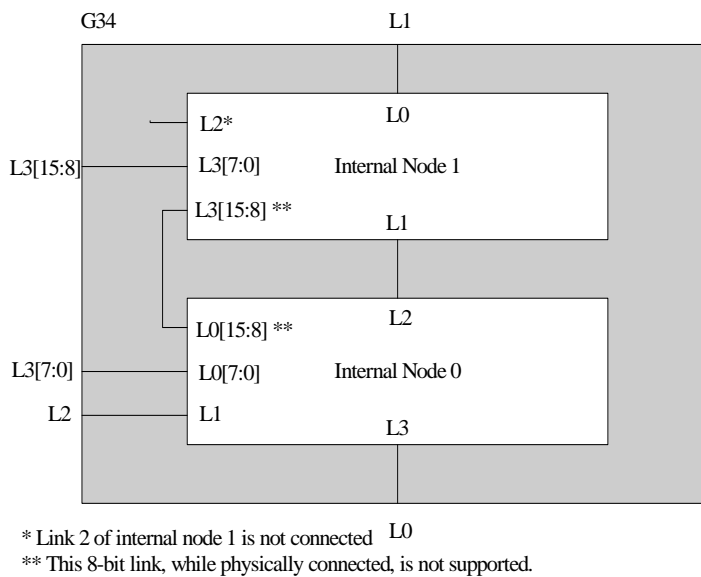


Figure 15: Dual-Node Processor Link Diagram

Since the internal 8 bit link is not supported, BIOS must do the following for dual-node processors:

For internal node 0 ([D18F3xE8\[IntNodeNum\]](#)==00b):

- If package link L3[7:0] is connected ([F0x98\[LinkCon\]](#)==1; see [D18F0x\[F8,D8,B8,98\]](#)):
 - Program [F0x84\[WidthIn, WidthOut\]](#)=000b. See [D18F0x\[E4,C4,A4,84\]](#).
 - Program [F0x170\[Ganged\]](#)=1. See [D18F0x\[18C:170\]](#).
- If package link L3[7:0] is not connected ([F0x98\[LinkCon\]](#)==0; see [D18F0x\[F8,D8,B8,98\]](#)):
 - Program [D18F0x16C\[ConnDly\]](#)=1.
 - Program [F4x84\[TransOff, EndOfChain\]](#)=1. See [D18F0x\[E4,C4,A4,84\]](#).

For internal node 1 ([D18F3xE8\[IntNodeNum\]](#)==01b):

- If package link L3[15:8] is connected ([F0xF8\[LinkCon\]](#)==1; see [D18F0x\[F8,D8,B8,98\]](#)):
 - Program [F0xE4\[WidthIn, WidthOut\]](#)=000b. See [D18F0x\[E4,C4,A4,84\]](#).
 - Program [F0x18C\[Ganged\]](#)=1. See [D18F0x\[18C:170\]](#).
- If package link L3[15:8] is not connected ([F0xF8\[LinkCon\]](#)==0; see [D18F0x\[F8,D8,B8,98\]](#)):
 - Program [D18F0x16C\[ConnDly\]](#)=1.
 - Program [F4xE4\[TransOff, EndOfChain\]](#)=1. See [D18F0x\[E4,C4,A4,84\]](#).

2.12.2 Termination and Compensation

The links are designed to operate in DC-termination mode as follows.




Figure 16: Link DC Termination Mode

Ron and Rtt are constructed with an array of parallel resistors that can be enabled or disabled to vary the resulting resistance. Each parallel resistor is referred to as a *tap*. Precision external resistors are used by the processor to determine the number of taps that must be enabled in order to match Ron and Rtt to the proper target values. The results of this compensation circuitry are observable in [D18F4x1\[9C,94,8C,84\]_xE0 \[Link Phy Compensation and Calibration Control 1\]\[RonRawCal\]](#) and [D18F4x1\[9C,94,8C,84\]_xE0 \[Link Phy Compensation and Calibration Control 1\]\[RttRawCal\]](#). Other fields in these registers are provided to offset the raw calculated compensation values or override them.

Compensation updates start after PWROK becomes valid (and occur while RESET_L is asserted).

The transmitter and receiver tristate in the PHY_OFF state, as entered by [D18F0x\[E4,C4,A4,84\]\[TransOff\]=1](#).

2.12.3 Equalization

A high speed data stream passing through the channel distorts due to various effects. The processor employs equalization to counter this problem and to improve electrical fidelity of the links. Equalization is employed by changing the voltage level transmitted before and after bit transitions. The transmitter can be attenuated to levels that vary based on bit history, as specified by [D18F4x1\[9C,94,8C,84\]_x\[7:6\]\[9:0\]\[8,0\]C_dm\[1\] \[Link Phy Tx Deemphasis and Margin Test Control\]](#). Equalization is not used at Gen1 frequencies.

2.12.4 Link Bandwidth Requirements

The bandwidth of a link may not exceed the bandwidth capacity of the nodes NB. The requirements are as follows:

- Where [NBCOF_MIN](#) is the NBCOF for the lowest frequency enabled NB P-state and HTCLK is the frequency of the link clock:
 - [D18F0x\[14C:130\]\[RetryModeEnable\]=1](#): $NBCOF_MIN \geq HTCLK / 2$.
 - [D18F0x\[14C:130\]\[RetryModeEnable\]=0](#): $NBCOF_MIN \geq HTCLK$.
- Both sides of the link must satisfy the following rules for all NB P-states:
 - $HTCLK > NBCOF$ for both sides of the link.
 - $HTCLK \leq NBCOF$ for both sides of the link.

2.12.5 Link Retry

The links support the error-retry mode described by the link specification, controlled by [D18F0x\[14C:130\] \[Link Retry\]](#) and [D18F0x150 \[Link Global Retry Control\]](#). Some requirements for operation in this mode:

- The processor does not support error-retry mode over links operating at Gen1 frequencies.
- IO links operating at Gen3 frequencies are required to have error-retry enabled. Coherent links operating

at Gen3 frequencies are expected to have error-retry enabled.

- Must be modified only when all links are operating at Gen1 frequency: [D18F0x150](#)[HtRetryCrcDatIns, HtRetryCrcDatInsDynEn, HtRetryCrcCmdPack, HtRetryCrcCmdPackDynEn].
- If any coherent links have error-retry enabled, then all coherent links are required to have it enabled after the coherent fabric has been configured. Until the warm reset to enable retry on all links, only configuration space cycles may be used in the coherent fabric.
- The retry history buffer for each ganged link supports up to 32 packets (each packet may include command and data), 16 packets for each unganged sublink.

2.12.6 Link LDTSTOP_L Disconnect-Reconnect

When disconnected for an LDTSTOP_L assertion, the state of the link and the reconnect time is a function of the link generation (Gen1 or Gen3) being used (or that the link is changing to, as a result of the LDTSTOP_L assertion), [D18F0x\[E4,C4,A4,84\]](#)[LdtStopTriEn], and [D18F0x\[18C:170\]](#)[LS2En] as follows:

Table 116: Link Disconnect Controls

Link Gen	LdtStopTriEn	LS2En	CLK	CAD, CTL	Reconnect delay
Gen1	0	X	L0 ¹	L0 ¹	Fast (about 1 us) ⁴
Gen1	1	0	L0 ¹	High imp ³	Fast (about 1 us) ⁴
Gen1	1	1	High imp ³	High imp ³	D18F3xD8 [ReConDel] ⁴
Gen3	X	0	L0 ¹	EI ²	D18F0x16C [T0Time]
Gen3	X	1	EI ²	EI ²	D18F0x16C [T0Time]

1. L0 represents the active, driven state.
 2. Electrical idle.
 3. High impedance.
 4. [D18F0x\[E4,C4,A4,84\]](#)[ExtCTL]=1 adds 50 us after CTL asserts.

2.12.7 LDTSTOP_L Requirements

- The processor requires additional minimum LDTSTOP_L assertion time for certain system configurations.
 - If any of the following system configuration properties are true the minimum LDTSTOP_L assertion time required by the processor is 10 us:
 - If [CPUID Fn8000_0001_EBX](#)[PkgType] is [C32r1](#) or [G34r1](#).
 - The system is in the C1E state, caches are not flushed, and [D18F3x\[84:80\]](#)[CpuPrbEn] for C1E is 0.
 - If [CPUID Fn8000_0001_EBX](#)[PkgType] is [AM3r2](#) then the minimum LDTSTOP_L assertion time required by the processor is 5 us.
- The processor requires additional STOP_GRANT to LDTSTOP_L deassertion time for certain system configurations.
 - LDTSTOP_L must not deassert less than 15 us after the processor broadcasts the STOP_GRANT message if [D18F3x\[84:80\]](#)[CpuPrbEn]=1 for the SMAF associated with the STPCLK assertion message and the link is operating at a Gen 1 frequency.
 - For all cases of LDTSTOP_L assertion, LDTSTOP_L must not deassert less than 14 us after the processor broadcasts the STOP_GRANT message.
 - For processors with an NB COF < 1.4 GHz, LDTSTOP_L must not deassert less than 20 us after the processor broadcasts the STOP_GRANT message.
 - If core power gating is enabled and the minimum time between power gating compute units is >= 5us and the link is operating at Gen 1 frequency then the minimum delay from STOP_GRANT to LDTSTOP_L deassertion must be increased. The condition when additional delay must be provided is

((D18F4x11[C:8][PwrGateEn] && (COUNT(D18F5x80[Enabled])>1) && (D18F5x128[PwrGateTmr]>=10b) && Gen1). The additional time required is: (COUNT(D18F5x80[Enabled])-1) * (UNIT(D18F5x128[PwrGateTmr] us)- 1 us). This time requirement is in addition to any other time requirements specified in this section.

- If core power gating is enabled and the link is operating at Gen 1 frequency then the minimum delay from STOP_GRANT to LDTSTOP_L deassertion must be increased. The time requirements for the following 2 cases are in addition to any other time requirements specified in this section.
 - If (D18F4x11[C:8][PwrGateEn] && D18F5x128[CC6PwrDwnRegEn] && (D18F5x128[PIIRegTime]>10b) && Gen1) then the additional time required is (COUNT(D18F5x80[Enabled])) * (UNIT(D18F5x128[PIIRegTime] us) - 1.5 us).
 - If (D18F4x11[C:8][PwrGateEn] && (D18F3xA0[PIILockTime]>001b) && Gen1) then the additional time required is (COUNT(D18F5x80[Enabled])) * (UNIT(D18F3xA0[PIILockTime] us) - 2 us).
- The processor requires a minimum LDTSTOP_L deassertion time as follows:
 - 3 us for unbuffered DIMMs and D18F2x90_dct[1:0][DisDIIShutdownSR]=1
 - 8 us for unbuffered DIMMs and D18F2x90_dct[1:0][DisDIIShutdownSR]=0.
 - 9 us for registered DIMMs.
- Narrow and slow links and use of D18F0x[E4,C4,A4,84][ExtCTL] can greatly increase the time for a Gen1 link to disconnect and reconnect, so the time between LDTSTOP_L assertions must be increased appropriately as required by 8.3 of the link specification, titled “Disconnecting and Reconnecting HyperTransport™ Links”.
- LDTSTOP_L may be asserted for the following reasons:
 - NB P-state transitions.
 - ACPI power state transitions controlled by D18F3x[84:80] (S4/S5, S3, S1, C1E, Link Init). See 2.6.1 [ACPI Power State Transitions].
 - In order to apply a change to a register field. This includes, but is not limited to:
 - Link width change. See D18F0x[E4,C4,A4,84][WidthOut, WidthIn].
 - Entry into ILM or BIST mode. See D18F0x[18C:170].
 - Link FIFO read pointer change. See D18F4x1[9C,94,8C,84]_x[DF,CF].
- LDTSTOP_L may not be asserted for the following reasons:
 - During AM3 0.2 GHz boot frequency.

2.12.8 Response Ordering

The processor supports non-standard response ordering, not required by the link specification. If the processor receives multiple IO-sourced memory read requests with certain attributes, then the processor ensures that the order of the responses to these requests is the same as the order in which the requests were received. The required attributes are:

- The requests have the same UnitID value (or logical UnitID if multiple UnitIDs are clumped; see D18F0x[11C,118,114,110] [Link Clumping Enable]).
- The requests have the same, non-zero SeqID value.
- The requests have the same PassPW bit value.
- The requests have the same Coherent (probe) bit value.
- The requests have the same RespPassPW bit value.
- The requests have the same Normal/Isochronous bit value.

This feature may allow IO devices to be designed that do not require re-order buffers. This behavior may be disabled through MSRC001_001F [NB Configuration 1 (NB_CFG1)][DisOrderRdRsp].

2.12.9 Link Testing, BIST, and ILM

The processor includes a link-defined BIST engine for each link. The control registers are found starting at [D18F4x1\[9C,94,8C,84\]_x100 \[Link BIST Control\]](#). See the link specification for more information.

The processor also supports link-defined internal loopback mode (ILM), controlled by [D18F0x\[18C:170\] \[Link Extended Control\]\[ILMEn\]](#).

2.12.10 Miscellaneous Behaviors and Requirements

- The processor does not support the link-defined Atomic read-modify-write command and returns target abort for any that are received.
- The processor does not support Device Messages and returns master abort for any that are received.
- The processor ignores the Chain bit.
- The processor checks for differential signaling on CTL[1:0] and disabled unused sublinks.
- Link width changes via LDTSTOP_L assertion are only supported by BIOS when the system is quiesced. Link frequency may only be changed by BIOS with a warm reset.
- Link frequency changes from one Gen3 frequency to another are not allowed.
- The processor register space does not include the Gen3 link-defined UCC bit or CPIC bit. However, functionally, the initial revisions of the processor would have these bits set to indicate that unthrottled command generation from IO links is supported (i.e., setting LinkTrain[DisCmdThrt] on the other side of the link) and command packet insertion from IO links is supported (i.e., setting LinkTrain[CPIEn] on the other side of the link). However, no assurances are made regarding future processor revisions; they may rely on throttling and disabled command packet insertion to operate.
- While transmitting to an IO link, the processor does not ever insert commands (other than NOPs) into data packets and the processor supports throttling command generation based on the state of [D18F0x168\[DisNcHtCmdThrottle\]](#).
- The processor logically supports link-defined mode combinations as follows (however electrical requirements may limit some options):

Table 117: Supported Link Operational Modes

Frequency	200-1000MHz	1200MHz and higher		
	Coupling/ Link Type	DC non-coherent operational	DC coherent operational	DC test/debug
Termination	RXDIF	RXDIF	RXDIF	RXDIF
Scrambling	No	Required	Required	Optional
Gen3 Training	No	Yes	Yes	Yes
Retry	No	Required	Optional at low speed	Optional

- The processor supports link-defined INTx messages. It emulates the ORing of INTx assertions throughout the system and broadcasts the result. To accomplish this, the processor uses separate counters for each of the four interrupts (INTA, INTB, INTC, and INTD) which track INTx assertions and deassertions received by the coherent fabric. Each assertion causes the counter to increment and each deassertion causes the counter to decrement. As each counter transitions from 0 to 1, the interrupt assertion message is broadcast. As each counter transitions from 1 to 0, the interrupt deassertion message is broadcast.
- The processor reflects system management messages E2h to FFh for vendor-defined virtual wire messages. Devices that send or receive them must have programmable registers to control the command encodings used so that different devices can interoperate.

- Ganged links leave the upper sublink driven after cold reset (per InLnSt) unless the lower sublink is unconnected. If the lower sublink of a ganged link is unconnected, the entire link is disabled.
- The processor cannot be used in a system where the sideband signal (RESET_L or LDTSTOP_L) skew between devices is greater than 100 us.
- The processor only supports synchronous clocking mode, where both sides of the link have their clocks derived from the same oscillator.
- No ordering may be assumed between broadcast requests and posted writes.
- The processor allows reordering of upstream posted requests with PassPW=1.

2.12.10.0.0.1 Legacy x87 FPU Stack Registers

.

2.13 RAS and Advanced Server Features

This section applies reliability, availability, and serviceability, or RAS, and related advanced server considerations.

2.13.1 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\[MCE\]](#).

2.13.1.1 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), including L3 cache and IO links. These MSRs 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 118 identifies the registers associated with each register bank:

Table 118: 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 MSRC000_0409	MSRC001_0048
MC5	MSR0000_0414	MSR0000_0415	MSR0000_0416	MSR0000_0417	MSRC001_0049
MC6	MSR0000_0418	MSR0000_0419	MSR0000_041A	MSR0000_041B	MSRC001_004A

Each register bank implements a minimum of 1 and a maximum of 9 machine check miscellaneous registers. A value of zero in $MCi_MISC0[BlkPtr]$ indicates that only one machine check miscellaneous register has been implemented for a given bank. If additional MISC registers beyond the first one are implemented for a given bank, a non-zero value in $MCi_MISC0[BlkPtr]$ points to the contiguous block of additional registers. The miscellaneous registers are denoted as MCi_MISCj , where j represents the number of miscellaneous registers for that bank.

The presence of valid information in the first MISC register in the bank (MCi_MISC0) is indicated by $MCi_STATUS[MiscV]$. The presence of valid information in additional implemented MISC registers is indicated by $MCi_MISCj[Val]$ in the target register.

2.13.1.2 Machine Check Errors Classes

The classes of machine check errors are:

- Uncorrectable
- Deferred
- Correctable

Uncorrectable errors cannot be corrected by hardware and may cause loss of data, corruption of processor state, or both. Uncorrectable errors update the status and address registers if not masked from logging in MCi_CTL_MASK . Information in the status and address registers from a previously logged deferred or correctable error is overwritten. Uncorrectable errors that are enabled in MCi_CTL result in reporting to software via machine check exceptions. 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, enable reporting of unmasked uncorrectable errors for normal operation. Disable reporting of uncorrectable errors only for debug purposes.

Deferred errors are errors that cannot be corrected by hardware, but do not cause an immediate interruption in program flow, loss of data integrity, or corruption of processor state. These errors indicate that data has been

corrupted but not consumed; no exception is generated because the data has not been referenced by a core or an IO link. Hardware writes information to the status and address registers in the corresponding bank that identifies the source of the error if deferred errors are enabled for logging. If there is information in the status and address registers from a previously logged correctable error, it is overwritten. Previously logged deferred or uncorrectable errors are not overwritten. Deferred errors are not reported via machine check exceptions; they can be seen by polling the MCI_STATUS registers. Refer to section 2.13.1.10 [Deferred Errors and Data Poisoning] for more detail on deferred errors.

Correctable errors are those which have been corrected by hardware 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). Hardware writes the status and address registers in the corresponding register bank with information that identifies the source of the error if they are enabled for logging. Correctable errors are not reported via machine check exceptions; some correctable errors may be reported to software via error thresholding (see 2.13.1.7 [Error Thresholding]).

The implications of these categories of errors are:

1. Uncorrected error; hardware did not deal with the problem.
 - Operationally (error handling), action required, 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.
 - Poisoned data consumed (as opposed to created), no way to avoid passing it to process or link.
2. Deferred error; hardware partially dealt with the problem via containment.
 - Operationally, action optional, 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; hardware dealt with the problem.
 - Operationally, no action required, 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 2.13.3 [Error Injection and Simulation] for more detail.

2.13.1.3 Error Detection, Action, Logging, and Reporting

Error detection is controlled by the MASK registers:

- Error detection for MCA controlled errors is enabled if not masked by MCI_CTL_MASK (see Table 118).
- Error masking is performed regardless of MCA bank enablement in MCG_CTL (MSR0000_017B).

Error action refers to the hardware response to an error, aside from logging and reporting. Enablement of error action for each error is enumerated in the EAC (Error Action Condition) column of the error descriptions tables as follows:

- D: Detected. The error action is taken if the error is detected (i.e., not masked). These actions occur regardless of whether the MCA bank is enabled in MCG_CTL.
- E: Enabled. The error action is taken if the the error is detected and the bank is enabled in MCG_CTL.

Error logging refers to the storing of information in the status registers, and is enabled if all of the following are true:

- Error detection is enabled.
- The MCA bank is enabled in MCG_CTL.

Error reporting refers to active notification of errors to software via machine check exceptions, and 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 uncorrectable.
- The error is enabled for reporting.
- CR4.MCE is enabled.

Notes:

1. If CR4.MCE is clear, an error configured to cause a machine check exception will cause a 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 2.13.1.7 [Error Thresholding]). Although no machine check exception will be generated, these notifications can be viewed as “correctable machine check interrupts”.

2.13.1.3.1 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 machine check reporting is disabled at the system level (CR4.MCE=0).
- Attempting to generate an MCE when a machine check is in progress on the same core (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 288 [EX Error Descriptions].
- EX “Mapper checkpoint array parity” error if UC=1. See Table 288 [EX Error Descriptions].

2.13.1.3.2 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\[Overflow\]](#).

Overflow alone does not indicate a shutdown condition. 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 [2.13.1.6 \[Handling Machine Check Exceptions\]](#)). If PCC is not indicated, any MCA data lost due to overflow was informational only and not critical to system hardware operation.

Table 119 indicates which errors are overwritten in the error status registers.

Table 119: Overwrite Priorities for All Banks

			Older Error					
			Uncorrectable		Deferred		Correctable	
			Enabled	Disabled	Enabled	Disabled	Enabled	Disabled
Newer Error	Uncorrectable	Enabled	-	-	Overwrite	Overwrite	Overwrite	Overwrite
		Disabled	-	-	Overwrite	Overwrite	Overwrite	Overwrite
	Deferred	Enabled	-	-	-	-	Overwrite	Overwrite
		Disabled	-	-	-	-	Overwrite	Overwrite
	Correctable	Enabled	-	-	-	-	-	-
		Disabled	-	-	-	-	-	-

2.13.1.4 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 by the operating system 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 118](#)). If initializing after a warm reset, then BIOS should check for valid MCA errors and if present save the status for later use (see [2.13.1.6 \[Handling Machine Check Exceptions\]](#)).

BIOS may initialize the 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). Alternatively, 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 2.13.1.3 [Error Detection, Action, Logging, and Reporting].)

Before ECC memory has been initialized with valid ECC check bits, BIOS must ensure that no memory operations are initiated if MCA reporting is enabled. This includes memory operations that may be initiated by hardware prefetching or other speculative execution. It is recommended that, until all of memory has been initialized with valid ECC check bits, the BIOS either does not have any valid MTRRs specifying a DRAM memory type or does not enable DRAM ECC machine check exceptions.

2.13.1.5 Error Code

The MCI_STATUS[ErrorCode] field contains information on the logged error. Table 120 [Error Code Types] identifies how to decode ErrorCode. The MCI_STATUS[ErrorCodeExt] field contains detailed, model-specific information that is used to further narrow identification for error diagnosis, but not error handling by software; see 2.13.1.6 [Handling Machine Check Exceptions].

For a given error reporting bank, Error Code is used in conjunction with the Extended Error Code (MCI_STATUS[ErrorCodeExt]) to uniquely identify the Error Type; the value of ErrorCodeExt is unique within Error Code Type. Details for each Error Type are described in the error signatures tables accompanying the MCI_STATUS register for each bank:

- MC0; Table 270.
- MC1; Table 273.
- MC2; Table 277.
- MC3; Reserved.
- MC4; Table 280 and Table 281.
- MC5; Table 289.
- MC6; Table 292.

Table 120: 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 121: Error Codes: Transaction Type

TT	Transaction Type
00	I: Instruction
01	D: Data
10	G: Generic
11	Reserved

Table 122: Error codes: cache level

LL	Cache Level
00	Reserved
01	L1: Level 1
10	L2: Level 2
11	LG: Generic (includes L3 cache)

Table 123: 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
0110	Prefetch
0111	Evict
1000	Probe (Snoop)

Table 124: 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 125: Error Codes: Memory or IO

II	Memory or IO
00	MEM: Memory Access

Table 125: Error Codes: Memory or IO

II	Memory or IO
01	Reserved
10	IO: IO Access
11	GEN: Generic

2.13.1.6 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 by software. 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 [2.13.1.6.1 \[MCA Differentiation Between System-Fatal and Process-Fatal Errors\]](#).

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 all 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 should be examined only when its valid bit is set.
 - 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 [2.13.1.5 \[Error Code\]](#) 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 sub-type 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 `MCi_STATUS` registers to see if error recovery is possible.
 - Error recovery is not possible when the processor context corrupt indicator (`MCi_STATUS[PCC]`) is set to 1.
 - The error overflow status indicator (`MCi_STATUS[Overflow]`) does not indicate whether error recovery is possible. See [2.13.1.3.2 \[Error Logging During Overflow\]](#).
 - If error recovery is not possible, the handler should log the error information and return to the operating system for system termination.
 - Check `MCi_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 [2.13.1.6.1 \[MCA Differentiation Between System-Fatal and Process-Fatal Errors\]](#) 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 `MCi_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.

2.13.1.6.1 MCA Differentiation Between System-Fatal and Process-Fatal Errors

Software may use MCA status to determine the scope errors and the recoverability of the system. The bits `MCi_STATUS[PCC]`, `MSR0000_017A[RIPV]`, and `MSR0000_017A[EIPV]` form a hierarchy. [Table 126](#) shows how these bits are interpreted.

Table 126: Error Scope Hierarchy

PCC	UC	RIPV	EIPV	Deferred	Poison	Comments
1	1	-	-	-	-	System fatal error. Signaled via machine check exception, action required. Error has corrupted system state (PCC=1). The error is fatal to the system and the system processing must be terminated.
0	1	1	1	-	-	Hardware uncorrectable, software restartable error. Signaled via machine check exception, action required. The error is confined to the process (PCC=0), and the process program can be restarted (RIPV=1) if the uncorrectable error is corrected by software.
0	1	0	-	-	0/1	Hardware uncorrectable, software containable error. Signaled via machine check exception, action required. The error is confined to the process, however the process cannot be restarted even if the uncorrectable error is corrected by software. Poison=1; the error is due to consumption of poison data. If the affected process or virtual machine is terminated, the system may continue operation.
0	0	-	-	1	0	Deferred error. Action optional. A latent error has been discovered, but not yet consumed; a machine check exception will be generated if the affected data is consumed. Error handling software may attempt to correct this data error, or prevent access by processes which map the data, or make the physical resource containing the data inaccessible. Note: May be detected on a demand access or a scrub access.
0	0	-	-	0	0	Correctable error. Signaled via error thresholding mechanisms (2.13.1.7 [Error Thresholding]); no action required.

2.13.1.7 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 signal 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, deferred, 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 270 \[LS Error Signatures\]](#).
- IF (MC1)
 - IF errors are counted and polled via [MSR0000_0407](#).
 - IF errors are listed in [Table 273 \[IF Error Signatures\]](#).
- CU (MC2)
 - CU errors are counted and polled via [MSR0000_040B](#).
 - CU errors are listed in [Table 277 \[CU Error Signatures\]](#).
- DRAM (MC4)
 - Memory errors are counted and polled or reported via [MSR0000_0413](#).
 - DRAM errors are the errors listed in [Table 279 \[NB Error Descriptions\]](#) as “D” (DRAM) in the ETG (Error Threshold Group) column.
 - Operating systems can avoid or stop using memory pages with excessive errors.
 - Spare memory can dynamically replace memory with excessive errors. See [2.10.9 \[On-Line Spare\]](#).
- Links (MC4)
 - Link errors are counted and polled or reported via [MSRC000_0408](#).
 - Link errors are the errors listed in [Table 279 \[NB Error Descriptions\]](#) 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 [2.12 \[Links\]](#) for details and restrictions on configuring links.
- L3 cache (MC4)
 - L3 cache errors are counted and polled or reported via [MSRC000_0409](#).
 - L3 cache errors are the errors listed in [Table 279 \[NB Error Descriptions\]](#) as “C” (Cache) in the ETG (Error Threshold Group) column.
 - If the product does not include an L3 cache, per [CPUID Fn8000_0006_EDX \[L3 Cache Identifiers\]](#), then the Valid and CntP bits are both 0 and the register logs no information.
- EX (MC5)
 - EX errors are counted and polled via [MSR0000_0417](#).
 - EX errors are listed in [Table 289 \[EX Error Signatures\]](#).

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.

2.13.1.8 Scrub Rate Considerations

This section gives guidelines for the scrub rate settings available in [D18F3x58 \[Scrub Rate Control\]](#). Scrubbers are used to periodically read cacheline sized data locations and associated tags. There are two primary benefits

to scrubbing. First, scrubbing corrects any correctable errors which are discovered before they can migrate into uncorrectable errors. This is particularly important for soft errors, which are caused by external sources such as radiation and which are temporary conditions which do not indicate malfunctioning hardware. Second, scrubbers help identify marginal or failed hardware by finding and logging repeated errors at the same locations (see also 2.13.1.7 [Error Thresholding]).

There are many factors which influence scrub rates. Among these are:

- The size of memory or cache to be scrubbed
- Resistance to upsets
- Geographic location and altitude
- Alpha particle contribution of packaging
- Performance sensitivity
- Risk aversion

The baseline recommendations in Table 127 are intended to provide excellent protection at most geographic locations, while having no measurable effect on performance. Adjustments may be necessary due to special circumstances. Refer to JEDEC standards for guidelines on adjusting for geographic location.

Table 127: Recommended Scrub Rates per Node

Register	Memory Size per Node (GB)	Register Setting	Scrub Rate
D18F3x58[L3Scrub]	-	10h	1.31 ms
D18F3x58[DramScrub]	0 GB == Size	00h	Disabled
	0 GB < Size <= 1 GB	12h	5.24 ms
	1 GB < Size <= 2 GB	11h	2.62 ms
	2 GB < Size <= 4 GB	10h	1.31 ms
	4 GB < Size <= 8 GB	0Fh	655.4 us
	8 GB < Size <= 16 GB	0Eh	327.7 us
	16 GB < Size <= 32 GB	0Dh	163.8 us
	32 GB < Size <= 64 GB	0Ch	81.9 us
	64 GB < Size <= 128 GB	0Bh	41.0 us
	128 GB < Size <= 256 GB	0Ah	20.5 us
	256 GB < Size	09h	10.2 us

For steady state operation, finding a range of useful scrub rates may be performed by selecting a scrub rate which is high enough to give good confidence about protection from accumulating errors and low enough that it has no measurable effect on performance. The above baselines are made to maximize error coverage without affecting performance and not based on specific processor soft error rates.

For low power states in which the processor core is halted, the power management configuration may affect scrubbing; see 2.9.5 [Memory Scrubbers].

2.13.1.9 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 DRAM address or a cache index. Logical addresses, while identifying the data, do not identify the location of the data.

If possible, physical storage locations in DRAM or 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.

2.13.1.9.1 Common Diagnosis Information

A common set of diagnosis information is useful for many problems. [Table 128](#) 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 128: 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 MSRC000_0409 D18F3x54 D18F3x58 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 D18F0x[E4,C4,A4,84][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 D18F0x[E4,C4,A4,84][CrcErr and CrcFloodEn].
- Link not present does not save status in MC4_STATUS. See D18F0x[E4,C4,A4,84][InitComplete].

Other registers may be needed depending on the specific error symptoms.

2.13.1.10 Deferred Errors and Data Poisoning

Deferred errors indicate error conditions which could not be corrected, but which require no action (i.e., action optional). Data poisoning marks data which has encountered an uncorrectable error, so that it can be tracked until it is consumed or discarded. Together, data poisoning and deferred errors provide an infrastructure for

reducing the severity of errors and the number of system outages for some classes of errors.

Processor cores create poison data and deferred errors as identified in the Error Signatures tables. When poison data or data with an uncorrectable ECC error is consumed, a machine check exception for an uncorrected error is signaled (MCI_STATUS[UC]). If the data is poison, MCI_STATUS[Poison] is also set. The NB converts any poison data sent from the core to a machine check exception with error type Compute Unit Data Error. To understand the cause of a machine check exception due to Compute Unit Data Error, examine the core MCA status registers for deferred errors.

The deferred error is logged in the MCA registers for diagnostic purposes at the time the error is discovered and the data is poisoned. This deferred error can help identify the source of the error. The deferred error is logged independently of any associated poison data machine check. For example, it is possible for a cache eviction to result in a deferred error associated with the cache, and a corresponding machine check exception from the NB which receives the data.

2.13.1.10.1 Software Error Handling

General software error handling for machine check exceptions is described in [2.13.1.6 \[Handling Machine Check Exceptions\]](#). Specific considerations for deferred errors are described here.

A deferred error is logged on the processor which detected the underlying error condition. The deferred error is primarily used for collecting diagnostic information on the cause of the error. When performing diagnosis due to deferred errors and poison data, it is important to distinguish between the physical address of the data in error and the physical location which caused the error. Physical address is used for tracking data usage. Physical location is important for determining the hardware in error.

2.13.2 DRAM Considerations for ECC

DRAM is protected by an error correcting code (ECC). There are two different error correcting codes supported by the memory controller. Both DRAM error correcting codes feature an ECC word formed by a symbol based code. The primary difference between the codes is the symbol size. The x4 code uses thirty-six 4-bit symbols to make a 144-bit ECC word made up of 128 data bits and 16 check bits, and the x8 code uses eighteen 8-bit symbols to make a 144-bit ECC word made up of 128 data bits and 16 check bits. BIOS must select one code by setting `D18F3x180[EccSymbolSize]`; see [Table 129 \[Configuration Specific Recommended ECC Symbol Size\]](#).

The x4 code is a single symbol correcting (SSC) and a double symbol detecting (DSD) code. This means the x4 code is able to correct 100% of single symbol errors (any bit error combination within one symbol), and detect 100% of double symbol errors (any bit error combination within two symbols).

The x8 code is a single symbol correcting code. The code is enhanced by an automated hardware history mechanism; see [2.13.2.3 \[Hardware Managed ECC History Mechanism\]](#) for more details. The x8 code can be further augmented with an optional software managed history mechanism to provide additional coverage against double symbol errors. The software mechanism extends the hardware history mechanism past a single cache line; see [2.13.2.4 \[Software Managed Bad Symbol Identification\]](#) for more details.

2.13.2.1 ECC Mapping

To understand the effects of configuration on DRAM device errors, the ECC must be mapped to the physical memory configuration, including device width, and DCT data interleaving.

- DRAM device width refers to the number of bits sourced simultaneously from a single memory chip. For example, a x4 device contributes 4-bits to the ECC word on each beat of data. In Figure 17, eighteen devices, representing one rank, contribute 72-bits on each beat.
- DCT data interleaving refers to the way bits from the two memory beats are organized to form an ECC word. When specified by D18F2x110[DctDatIntLv], even and odd bits from the two 72-bit beats can be interleaved to create the 144-bit ECC word as shown in Figure 17.

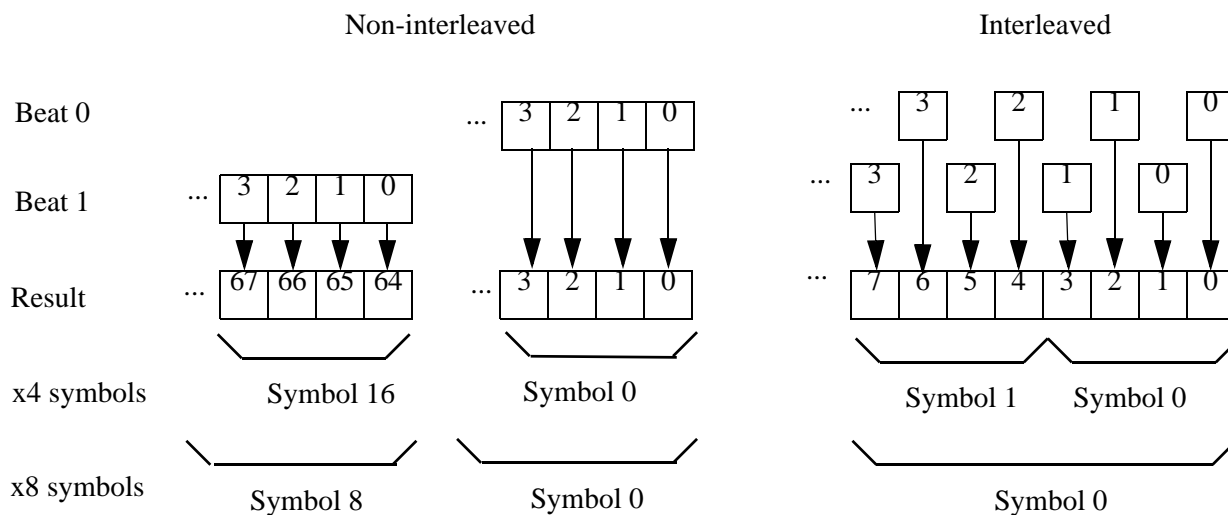


Figure 17: Example of Line Interleaving from x4 DRAM Using x4 and x8 ECC

When DCT data interleaving is enabled, a partially failing device (e.g., pin failure) contributes two incorrect bits to the same symbol. When DCT data interleaving is not enabled, a partially failing device contributes failures to two different symbols.

2.13.2.2 Single Device Data Correction (Chipkill)

In certain configurations, the ECC provides single DRAM device data correction, known as “chipkill” functionality; all single symbol errors caused by a failed DRAM device are corrected. Chipkill recovery is only possible when indicated by D18F3x44[ChipKillEccCap] and all bits within a symbol are sourced from a single DRAM device. For chipkill functionality, the symbol size must be at least twice the DRAM device width, since a given device contributes data to two DRAM beats, and DCT data interleaving must be enabled. For example, for a system configured with x4 devices, x8 symbols are necessary so that the failure of one device can be corrected since the device contributes 8 total bits of data confined to one symbol.

For configuration specific recommendations and detection and correction characteristics see Table 129 below.

Table 129: Configuration Specific Recommended ECC Symbol Size

DRAM Device Width	Recommended Symbol Size (D18F3x180[Ecc-SymbolSize])	Error Description
x4	x8	The failure of a DRAM device results in errors to one symbol, and can be corrected (chipkill).
x8 or greater	x8	The failure of a DRAM device results in errors to multiple symbols, and can be detected with a high probability but cannot be corrected. These configurations are not recommended for high reliability or high availability systems, due to the higher potential for uncorrected, undetected, or mis-corrected errors.

2.13.2.3 Hardware Managed ECC History Mechanism

The x8 code is enhanced by the following automated hardware history mechanism to provide additional coverage against two symbol errors that result from one or more defective DRAM devices:

- The hardware error history is cleared at the beginning of each cache line received from DRAM.
- For each ECC word in the cache line:
 - If the ECC indicates that there is no error, then the data can safely be forwarded.
 - If the ECC indicates that there is an uncorrectable error, then the ECC word is considered uncorrectable.
 - If the ECC indicates a correctable error, then:
 - If it is the first error in the cache line, the symbol position in error is saved for the remainder of the cache line (only the first error of the cache line is saved), and the data is corrected but not yet forwarded.
 - If the ECC indicates a subsequent correctable error in the same symbol position, then the data is corrected, but not yet forwarded.
 - If the ECC indicates subsequent correctable errors in symbols other than the saved symbol position, then all ECC words in the cache line are considered uncorrectable.
 - Uncorrectable errors detected by this mechanism will set both UECC=1 and CECC=1 in [MSR0000_0411](#).
- At the end of the cache line, all data is forwarded.

2.13.2.4 Software Managed Bad Symbol Identification

x8 ECC can be further augmented by a software managed algorithm to provide coverage against two symbol errors in the same ECC word by using error information across cache lines.

When software determines that a DRAM device is bad, it should program the rank and symbol into the BadDramCs and BadDramSymbol fields of the appropriate register, [D18F3x138 \[DCT0 Bad Symbol Identification\]](#) or [D18F3x13C \[DCT1 Bad Symbol Identification\]](#). The bad symbol is determined by using the syndrome captured during correctable errors, as described in section [2.13.2.5.1 \[x8 ECC\]](#). Determining that a device is bad can be performed using the guidelines in section [2.13.1.7 \[Error Thresholding\]](#). Hardware uses the information that a symbol is bad to protect against double symbol errors with the following history scheme:

- If an ECC word has a correctable error only in a rank and symbol identified by this register, then the error is a correctable error.
- If an ECC word has an error in a rank identified by this register, but a different symbol, then the word is

presumed to have at least two symbol errors and the error is treated as uncorrectable.

2.13.2.5 ECC Syndromes

For memory errors, the sections below describe how to find the DIMM in error. The process varies slightly according to the ECC code in use. To determine which ECC code is being used, see [D18F3x180\[EccSymbolSize\]](#).

For correctable errors, the DIMM in error is uniquely identified by the error address ([D18F3x54\[ErrAddr\]](#)) and the ECC syndrome ([D18F3x48\[Syndrome\[15:8\]\]](#) and [D18F3x4C\[Syndrome\[7:0\]\]](#)).

2.13.2.5.1 x8 ECC

The use of x8 ECC is indicated in [D18F3x180\[EccSymbolSize\]](#).

The syndrome field uniquely identifies the failing bit positions of a correctable ECC error. Only syndromes identified in [Table 130](#) are correctable by the error correcting code.

Symbols 00h-07h map to data bits 0-63; symbols 08h-0Fh map to data bits 64-127; symbol 10h maps to special purpose non-data bits; symbols 11h-12h map to ECC check bits for data bits 0-127 and symbol 10h.

To use [Table 130](#), find the 16-bit syndrome value in the table. Because of the large size of the table, this is performed by taking the low order byte of the syndrome as the row number, then scanning the row for the complete 16-bit syndrome. If it is not found, use the high order byte of the syndrome as the row number, then scan the row for the complete 16-bit syndrome. Once the syndrome is found in the table, the corresponding Symbol In Error column indicates which symbol, and therefore which DIMM has the error. The Error Bitmask column indicates the bits in error in the symbol.

For example, assume the ECC syndrome is 03EAh. First search row EAh for the complete syndrome. Since it is not found, search row 03h for the complete syndrome. It is found in column 9h, so symbol 9h has the error. Since the error bitmask indicates value 3h (0011b), bits 0 and 1 within that symbol are corrupted. Symbol 9h maps to bits 72-79, so the corrupted bits are 72 and 73 of the line.

Table 130: x8 ECC Correctable Syndromes

Error Bit-mask	Symbol In Error																		
	12h	11h	10h	Fh	Eh	Dh	Ch	Bh	Ah	9h	8h	7h	6h	5h	4h	3h	2h	1h	0h
00	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000	0000
01	0100	0001	0101	01B8	015C	012E	01C6	0163	01FD	0189	019D	B801	5C01	2E01	C601	6301	FD01	8901	9D01
02	0200	0002	0202	0201	02B8	025C	02FD	02C6	028B	0263	024B	0102	B802	5C02	FD02	C602	8B02	6302	4B02
03	0300	0003	0303	03B9	03E4	0372	033B	03A5	0376	03EA	03D6	B903	E403	7203	3B03	A503	7603	EA03	D603
04	0400	0004	0404	0402	0401	04B8	048B	04FD	0467	04C6	0496	0204	0104	B804	8B04	FD04	6704	C604	9604
05	0500	0005	0505	05BA	055D	0596	054D	059E	059A	054F	050B	BA05	5D05	9605	4D05	9E05	9A05	4F05	0B05
06	0600	0006	0606	0603	06B9	06E4	0676	063B	06EC	06A5	06DD	0306	B906	E406	7606	3B06	EC06	A506	DD06
07	0700	0007	0707	07BB	07E5	07CA	07B0	0758	0711	072C	0740	BB07	E507	CA07	B007	5807	1107	2C07	4007
08	0800	0008	0808	0804	0802	0801	0867	088B	08CE	08FD	085D	0408	0208	0108	6708	8B08	CE08	FD08	5D08
09	0900	0009	0909	09BC	095E	092F	09A1	09E8	0933	0974	09C0	BC09	E09	2F09	A109	E809	3309	7409	C009
0A	0A00	000A	0A0A	0A05	0ABA	0A5D	0A9A	0A4D	0A45	0A9E	0A16	050A	BA0A	5D0A	9A0A	4D0A	450A	9E0A	160A
0B	0B00	000B	0B0B	0BBD	0BE6	0B73	0B5C	0B2E	0BB8	0B17	0B8B	BD0B	E60B	730B	5C0B	2E0B	B80B	170B	8B0B
0C	0C00	000C	0C0C	0C06	0C03	0CB9	0CEC	0C76	0CA9	0C3B	0CCB	060C	030C	B90C	EC0C	760C	A90C	3B0C	CB0C
0D	0D00	000D	0D0D	0DBE	0D5F	0D97	0D2A	0D15	0D54	0DB2	0D56	BE0D	5F0D	970D	2A0D	150D	540D	B20D	560D
0E	0E00	000E	0E0E	0E07	0EBB	0EE5	0E11	0EB0	0E22	0E58	0E80	070E	BB0E	E50E	110E	B00E	220E	580E	800E
0F	0F00	000F	0F0F	0FBF	0FE7	0FCB	0FD7	0FD3	0FDF	0FD1	0F1D	BF0F	E70F	CB0F	D70F	D30F	DF0F	D10F	1D0F
10	1000	0010	1010	1008	1004	1002	10CE	1067	10ED	108B	10BA	0810	0410	0210	CE10	6710	ED10	8B10	BA10
11	1100	0011	1111	11B0	1158	112C	1108	1104	1110	1102	1127	B011	5811	2C11	0811	0411	1011	0211	2711
12	1200	0012	1212	1209	12BC	125E	1233	12A1	1266	12E8	12F1	0912	BC12	5E12	3312	A112	6612	E812	F112
13	1300	0013	1313	13B1	13E0	1370	13F5	13C2	139B	1361	136C	B113	E013	7013	F513	C213	9B13	6113	6C13
14	1400	0014	1414	140A	1405	14BA	1445	149A	148A	144D	142C	0A14	0514	BA14	4514	9A14	8A14	4D14	2C14

Table 130: x8 ECC Correctable Syndromes

Error Bit-mask	Symbol In Error																			
	12h	11h	10h	Fh	Eh	Dh	Ch	Bh	Ah	9h	8h	7h	6h	5h	4h	3h	2h	1h	0h	
15	1500	0015	1515	15B2	1559	1594	1583	15F9	1577	15C4	15B1	B215	5915	9415	8315	F915	7715	C415	B115	
16	1600	0016	1616	160B	16BD	16E6	16B8	165C	1601	162E	1667	0B16	BD16	E616	B816	5C16	0116	2E16	6716	
17	1700	0017	1717	17B3	17E1	17C8	177E	173F	17FC	17A7	17FA	B317	E117	C817	7E17	3F17	FC17	A717	FA17	
18	1800	0018	1818	180C	1806	1803	18A9	18EC	1823	1876	18E7	0C18	0618	0318	A918	EC18	2318	7618	E718	
19	1900	0019	1919	19B4	195A	192D	196F	198F	19DE	19FF	197A	B419	5A19	2D19	6F19	8F19	DE19	FF19	7A19	
1A	1A00	001A	1A1A	1A0D	1ABE	1A5F	1A54	1A2A	1AA8	1A15	1AAC	0D1A	BE1A	5F1A	541A	2A1A	A81A	151A	AC1A	
1B	1B00	001B	1B1B	1BB5	1BE2	1B71	1B92	1B49	1B55	1B9C	1B31	B51B	E21B	711B	921B	491B	551B	9C1B	311B	
1C	1C00	001C	1C1C	1C0E	1C07	1CBB	1C22	1C11	1C44	1CB0	1C71	0E1C	071C	BB1C	221C	111C	441C	B01C	711C	
1D	1D00	001D	1D1D	1DB6	1D5B	1D95	1DE4	1D72	1DB9	1D39	1DEC	B61D	5B1D	951D	E41D	721D	B91D	391D	EC1D	
1E	1E00	001E	1E1E	1E0F	1EBF	1EE7	1EDF	1ED7	1ECF	1ED3	1E3A	0F1E	BF1E	E71E	DF1E	D71E	CF1E	D31E	3A1E	
1F	1F00	001F	1F1F	1FB7	1FE3	1FC9	1F19	1FB4	1F32	1F5A	1FA7	B71F	E31F	C91F	191F	B41F	321F	5A1F	A71F	
20	2000	0020	2020	2010	2008	2004	20ED	20CE	20AB	2067	2005	1020	0820	0420	ED20	CE20	AB20	6720	0520	
21	2100	0021	2121	21A8	2154	212A	212B	21AD	2156	21EE	2198	A821	5421	2A21	2B21	AD21	5621	EE21	9821	
22	2200	0022	2222	2211	22B0	2258	2210	2208	2220	2204	224E	1122	B022	5822	1022	AD22	2022	0422	4E22	
23	2300	0023	2323	23A9	23EC	2376	23D6	236B	23DD	238D	23D3	A923	EC23	7623	D623	6B23	DD23	8D23	D323	
24	2400	0024	2424	2412	2409	24BC	2466	2433	24CC	24A1	2493	1224	0924	BC24	6624	3324	CC24	A124	9324	
25	2500	0025	2525	25AA	2555	2592	25A0	2550	2531	2528	250E	AA25	5525	9225	A025	5525	3125	2825	0E25	
26	2600	0026	2626	2613	26B1	26E0	269B	26F5	2647	26C2	26D8	1326	B126	E026	9B26	F526	4726	C226	D826	
27	2700	0027	2727	27AB	27ED	27CE	275D	2796	27BA	274B	2745	AB27	ED27	CE27	5D27	9627	BA27	4B27	4527	
28	2800	0028	2828	2814	280A	2805	288A	2845	2865	289A	2858	1428	0A28	0528	8A28	4528	6528	9A28	5828	
29	2900	0029	2929	29AC	2956	292B	294C	2926	2998	2913	29C5	AC29	5629	2B29	4C29	2629	9829	1329	C529	
2A	2A00	002A	2A2A	2A15	2AB2	2A59	2A77	2A83	2AEE	2AF9	2A13	152A	B22A	592A	772A	832A	EE2A	F92A	132A	
2B	2B00	002B	2B2B	2BAD	2BEE	2B77	2BB1	2BE0	2B13	2B70	2B8E	AD2B	EE2B	772B	B12B	E02B	132B	702B	8E2B	
2C	2C00	002C	2C2C	2C16	2C0B	2CBD	2C01	2CB8	2C02	2C5C	2CCE	162C	0B2C	BD2C	012C	B82C	022C	5C2C	CE2C	
2D	2D00	002D	2D2D	2DAE	2D57	2D93	2DC7	2DDB	2DFF	2DD5	2D53	AE2D	572D	932D	C72D	DB2D	FF2D	D52D	532D	
2E	2E00	002E	2E2E	2E17	2EB3	2EE1	2EFC	2E7E	2E89	2E3F	2E85	172E	B32E	E12E	FC2E	F22E	AA2E	3F2E	5E2E	
2F	2F00	002F	2F2F	2FAF	2FEF	2FCF	2F3A	2F1D	2F74	2FB6	2F18	AF2F	EF2F	CF2F	3A2F	1D2F	742F	B62F	182F	
30	3000	0030	3030	3018	300C	3006	3023	30A9	3046	30EC	30BF	1830	0C30	0630	2330	A930	4630	EC30	BF30	
31	3100	0031	3131	31A0	3150	3128	31E5	31CA	31B6	3165	3122	A031	5031	2831	E531	CA31	BB31	6531	2231	
32	3200	0032	3232	3219	32B4	325A	32DE	326F	32CD	328F	32F4	1932	B432	5A32	DE32	6F32	AA32	8F32	F432	
33	3300	0033	3333	33A1	33E8	3374	3318	330C	3330	3306	3369	A133	E833	7433	1833	0C33	3033	0633	6933	
34	3400	0034	3434	341A	340D	34BE	34A8	3454	3421	342A	3429	1A34	0D34	BE34	A834	5434	2134	2A34	2934	
35	3500	0035	3535	35A2	3551	3590	356E	3537	35DC	35A3	35B4	A235	5135	9035	6E35	3735	DC35	A335	B435	
36	3600	0036	3636	361B	36B5	36E2	3655	3692	36AA	3649	3662	1B36	B536	E236	5536	9236	AA36	4936	6236	
37	3700	0037	3737	37A3	37E9	37CC	3793	37F1	3757	37C0	37FF	A337	E937	CC37	9337	F137	5737	C037	FF37	
38	3800	0038	3838	381C	380E	3807	3844	3822	3888	3811	38E2	1C38	0E38	0738	4438	2238	8838	1138	E238	
39	3900	0039	3939	39A4	3952	3929	3982	3941	3975	3998	397F	A439	5239	2939	8239	4139	7539	9839	7F39	
3A	3A00	003A	3A3A	3A1D	3AB6	3A5B	3AB9	3AE4	3A03	3A72	3AA9	1D3A	B63A	5B3A	B93A	E43A	033A	723A	A93A	
3B	3B00	003B	3B3B	3BA5	3BEA	3B75	3B7F	3B87	3BFE	3BFB	3B34	A53B	EA3B	753B	7F3B	873B	FE3B	FB3B	343B	
3C	3C00	003C	3C3C	3C1E	3C0F	3CBF	3CCF	3CDF	3CEF	3CD7	3C74	1E3C	0F3C	BF3C	CF3C	DF3C	EF3C	D73C	743C	
3D	3D00	003D	3D3D	3DA6	3D53	3D91	3D09	3DBC	3D12	3D5E	3DE9	A63D	533D	913D	093D	BC3D	123D	5E3D	E93D	
3E	3E00	003E	3E3E	3E1F	3EB7	3EE3	3E32	3E19	3E64	3EB4	3EF3	1F3E	B73E	E33E	323E	193E	643E	B43E	3F3E	
3F	3F00	003F	3F3F	3FA7	3FEB	3FCD	3FF4	3F7A	3F99	3F3D	3FA2	A73F	EB3F	CD3F	F43F	7A3F	993F	3D3F	A23F	
40	4000	0040	4040	4020	4010	4008	40AB	40ED	4027	40CE	400A	2040	1040	0840	AB40	ED40	2740	CE40	0A40	
41	4100	0041	4141	4198	414C	4126	416D	418E	41DA	4147	4197	9841	4C41	2641	6D41	8E41	DA41	4741	9741	
42	4200	0042	4242	4221	42A8	4254	4256	422B	42AC	42AD	4241	2142	A842	5442	5642	2B42	AC42	AD42	4142	
43	4300	0043	4343	4399	43F4	437A	4390	4348	4351	4324	43DC	9943	F443	7A43	9043	4843	5143	2443	DC43	
44	4400	0044	4444	4422	4411	44B0	4420	4410	4440	4408	449C	2244	1144	B044	2044	1044	4044	0844	9C44	
45	4500	0045	4545	459A	454D	459E	45E6	4573	45BD	4581	4501	9A45	4D45	9E45	E645	7345	BD45	8145	0145	
46	4600	0046	4646	4623	46A9	46EC	46DD	46D6	46CB	466B	46D7	2346	A946	EC46	DD46	D646	6B46	D746		
47	4700	0047	4747	479B	47F5	47C2	471B	47B5	4736	47E2	474A	9B47	F547	C247	1B47	B547	3647	E247	4A47	
48	4800	0048	4848	4824	4812	4809	48CC	4866	48E9	4833	4857	2448	1248	0948	CC48	6648	E948	3348	5748	
49	4900	0049	4949	499C	494E	4927	490A	4905	4914	49BA	49CA	9C49	E49	2749	0A49	0549	1449	BA49	CA49	
4A	4A00	004A	4A4A	4A25	4AAA	4A55	4A31	4AA0	4A62	4A50	4A1C	254A	AA4A	554A	314A	A04A	624A	504A	1C4A	
4B	4B00	004B	4B4B	4B9D	4BF6	4B7B	4BF7	4BC3	4B9F	4BD9	4B81	9D4B	F64B	7B4B	F74B	C34B	9F4B	D94B	814B	
4C	4C00	004C	4C4C	4C26	4C13	4CB1	4C47	4C9B	4C8E	4CF5	4CC1	264C	134C	B14C	474C	9B4C	8E4C	F54C	C14C	
4D	4D00	004D	4D4D	4D9E	4D4F	4D9F	4D81	4DF8	4D73	4D7C	4D5C	9E4D	4F4D	9F4D	814D	F84D	734D	7C4D	5C4D	
4E	4E00	004E	4E4E	4E27	4EAB	4EED	4EBA	4E5D	4E05	4E96	4E8A	274E	AB4E	ED4E	BA4E	5D4E	054E	964E	8A4E	
4F	4F00	004F	4F4F	4F9F	4FF7	4FC3	4F7C	4F3E	4FF8	4F1F	4F17	9F4F	F74F	C34F	7C4F	3E4F	F84F	1F4F	174F	
50	5000	0050	5050	5028	5014	500A	5065	508A	50CA	5045	50B0	2850	1450	0A50	6550	8A50	CA50	4550	B050	
51	5100	0051	5151	5190	5148	5124	51A3	51E9	5137	51CC	512D	9051	4851	2451	A351	E951	3751	CC51	2D51	
52	5200	0052	5252	5229	52AC	5256	5298	524C	5241	5226	52FB	2952	AC52	5652	9852	4C52	4152	2652	FB52	
53	5300	0053	5353	5391	53F0	5378	535E	532F	53BC	53AF	5366	9153	F053	7853	5E53	2F53	BC53	AF53	6653	
54	5400	0054	5454	542A	5415	54B2	54EE	5477	54AD	5483	5426	2A54	1554	B254	EE54	7754	AD54	8354	2654	
55	5500	0055	5555	5592	5549	559C	5528	5514	5550	550A	55BB	9255	4955	9C55	2855	1455	5055	0A55	BB55	
56	5600	0056	5656	562B	56AD	56EE	5613	56B1	5626	56E0	566D	2B56	AD56	EE56	1356	B156	2656	E056	6D56	
57	5700	0057	5757	5793	57F1	57C0	57D5	57D2	57DB	5769	57F0	9357	F157	C057	D557	D257	DB57	6957	F057	

Table 130: x8 ECC Correctable Syndromes

Error Bit-mask	Symbol In Error																			
	12h	11h	10h	Fh	Eh	Dh	Ch	Bh	Ah	9h	8h	7h	6h	5h	4h	3h	2h	1h	0h	
58	5800	0058	5858	582C	5816	580B	5802	5801	5804	58B8	58ED	2C58	1658	0B58	0258	0158	0458	B858	ED58	
59	5900	0059	5959	5994	594A	5925	59C4	5962	59F9	5931	5970	9459	4A59	2559	C459	6259	F959	3159	7059	
5A	5A00	005A	5A5A	5A2D	5AAE	5A57	5AFF	5AC7	5A8F	5ADB	5AA6	2D5A	AE5A	575A	FF5A	C75A	8F5A	DB5A	A65A	
5B	5B00	005B	5B5B	5B95	5BF2	5B79	5B39	5BA4	5B72	5B52	5B3B	955B	F25B	795B	395B	A45B	725B	525B	3B5B	
5C	5C00	005C	5C5C	5C2E	5C17	5CB3	5C89	5CFC	5C63	5C7E	5C7B	2E5C	175C	B35C	895C	FC5C	635C	7E5C	7B5C	
5D	5D00	005D	5D5D	5D96	5D4B	5D9D	5D4F	5D9F	5D9E	5DF7	5DE6	965D	4B5D	9D5D	4F5D	9F5D	9E5D	F75D	E65D	
5E	5E00	005E	5E5E	5E2F	5EAF	5EEF	5E74	5E3A	5EE8	5E1D	5E30	2F5E	AF5E	EF5E	745E	3A5E	E85E	1D5E	305E	
5F	5F00	005F	5F5F	5F97	5FF3	5FC1	5FB2	5F59	5F15	5F94	5FAD	975F	F35F	C15F	B25F	595F	155F	945F	AD5F	
60	6000	0060	6060	6030	6018	600C	6046	6023	608C	60A9	600F	3060	1860	0C60	4660	2360	8C60	A960	0F60	
61	6100	0061	6161	6188	6144	6122	6180	6140	6171	6120	6192	8861	4461	2261	8061	4061	7161	2061	9261	
62	6200	0062	6262	6231	62A0	6250	62BB	62E5	6207	62CA	6244	3162	A062	5062	BB62	E562	0762	CA62	4462	
63	6300	0063	6363	6389	63FC	637E	637D	6386	63FA	6343	63D9	8963	FC63	7E63	7D63	8663	FA63	4363	D963	
64	6400	0064	6464	6432	6419	64B4	64CD	64DE	64EB	646F	6499	3264	1964	B464	CD64	DE64	EB64	6F64	9964	
65	6500	0065	6565	658A	6545	659A	650B	65BD	6516	65E6	6504	8A65	4565	9A65	0B65	3665	D865	1B65	0465	
66	6600	0066	6666	6633	66A1	66E8	6630	6618	6660	660C	66D2	3366	A166	E866	3066	1866	6066	0C66	D266	
67	6700	0067	6767	678B	67FD	67C6	67F6	677B	679D	6785	674F	8B67	FD67	C667	F667	7B67	9D67	8567	4F67	
68	6800	0068	6868	6834	681A	680D	6821	68A8	6842	6854	6852	3468	1A68	0D68	2168	A868	4268	5468	5268	
69	6900	0069	6969	698C	6946	6923	69E7	69CB	69BF	69DD	69CF	8C69	4669	2369	E769	CB69	BF69	DD69	CF69	
6A	6A00	006A	6A6A	6A35	6AA2	6A51	6ADC	6A6E	6AC9	6A37	6A19	356A	A26A	516A	DC6A	6E6A	C96A	376A	196A	
6B	6B00	006B	6B6B	6B8D	6BFE	6B7F	6B1A	6B0D	6B34	6BBE	6B84	8D6B	FE6B	7F6B	1A6B	0D6B	346B	BE6B	846B	
6C	6C00	006C	6C6C	6C36	6C1B	6CB5	6CAA	6C55	6C25	6C92	6CC4	366C	1B6C	B56C	AA6C	556C	256C	926C	C46C	
6D	6D00	006D	6D6D	6D8E	6D47	6D9B	6D6C	6DD8	6D1B	6D59	8E6D	476D	9B6D	6C6D	366D	D86D	1B6D	596D	566D	
6E	6E00	006E	6E6E	6E37	6EA3	6EE9	6E57	6E93	6EAE	6EF1	6E8F	376E	A36E	E96E	576E	936E	AE6E	F16E	8F6E	
6F	6F00	006F	6F6F	6F8F	6FFF	6FC7	6F91	6FF0	6F53	6F78	6F12	8F6F	FF6F	C76F	916F	F06F	536F	786F	126F	
70	7000	0070	7070	7038	701C	700E	7088	7044	7061	7022	70B5	3870	1C70	0E70	8870	4470	6170	2270	B570	
71	7100	0071	7171	7180	7140	7120	714E	7127	719C	71AB	7128	8071	4071	2071	4E71	2771	9C71	AB71	2871	
72	7200	0072	7272	7239	72A4	7252	7275	7282	72EA	7241	72FE	3972	A472	5272	7572	8272	EA72	4172	FE72	
73	7300	0073	7373	7381	73F8	737C	73B3	73E1	7317	73C8	7363	8173	F873	7C73	B373	E173	1773	C873	6373	
74	7400	0074	7474	743A	741D	74B6	7403	74B9	7406	74E4	7423	3A74	1D74	B674	0374	DA74	0674	E474	2374	
75	7500	0075	7575	7582	7541	7598	75C5	75DA	75FB	756D	75BE	8275	4175	9875	C575	DA75	FB75	6D75	BE75	
76	7600	0076	7676	763B	76A5	76EA	76FE	767F	768D	7687	7668	3B76	A576	EA76	FE76	7F76	8D76	8776	6876	
77	7700	0077	7777	7783	77F9	77C4	7738	771C	7770	770E	77F5	8377	F977	C477	3877	1C77	7077	0E77	F577	
78	7800	0078	7878	783C	781E	780F	78EF	78CF	78AF	78DF	78E8	3C78	1E78	0F78	EF78	CF78	AF78	DF78	E878	
79	7900	0079	7979	7984	7942	7921	7929	79AC	7952	7956	7975	8479	4279	2179	2979	AC79	5279	5679	7579	
7A	7A00	007A	7A7A	7A3D	7AA6	7A53	7A12	7A09	7A24	7ABC	7AA3	3D7A	A67A	537A	127A	097A	247A	BC7A	A37A	
7B	7B00	007B	7B7B	7B85	7BFA	7B7D	7BD4	7B6A	7BD9	7B35	7B3E	857B	FA7B	7D7B	D47B	6A7B	D97B	357B	7B7B	
7C	7C00	007C	7C7C	7C3E	7C1F	7CB7	7C64	7C32	7CC8	7C19	7C7E	3E7C	1F7C	B77C	647C	327C	C87C	197C	7E7C	
7D	7D00	007D	7D7D	7D86	7D43	7D9B	7DA2	7D51	7D35	7D90	7DE3	867D	437D	997D	A27D	517D	357D	907D	E37D	
7E	7E00	007E	7E7E	7E3F	7EA7	7EEB	7E99	7EF4	7E43	7E7A	7E35	3F7E	A77E	EB7E	997E	F47E	437E	7A7E	357E	
7F	7F00	007F	7F7F	7F87	7FFB	7FC5	7F5F	7F97	7FBE	7FF3	7FA8	877F	FB7F	C57F	5F7F	977F	BE7F	F37F	A87F	
80	8000	0080	8080	8040	8020	8010	8027	80AB	804E	80ED	8014	4080	2080	1080	2780	AB80	4E80	ED80	1480	
81	8100	0081	8181	81F8	817C	813E	81E1	81C8	81B3	8164	8189	F881	7C81	3E81	E181	C881	B381	6481	8981	
82	8200	0082	8282	8241	8298	824C	82DA	826D	82C5	828E	825F	4182	9882	4C82	DA82	6D82	C582	8E82	5F82	
83	8300	0083	8383	83F9	83C4	8362	831C	830E	8338	8307	83C2	F983	C483	6283	1C83	0E83	3883	0783	C283	
84	8400	0084	8484	8442	8421	84A8	84AC	8456	8429	842B	8482	4284	2184	A884	AC84	5684	2984	2B84	8284	
85	8500	0085	8585	85FA	857D	8586	856A	8535	85D4	85A2	851F	FA85	7D85	8685	3585	D485	A285	1F85	8585	
86	8600	0086	8686	8643	8699	86F4	8651	8690	86A2	8648	86C9	4386	9986	F486	5186	9086	A286	4886	C986	
87	8700	0087	8787	87FB	87C5	87DA	8797	87F3	875F	87C1	8754	FB87	C587	DA87	9787	F387	5F87	C187	5487	
88	8800	0088	8888	8844	8822	8811	8840	8820	8880	8810	8849	4488	2288	1188	4088	2088	8088	1088	4988	
89	8900	0089	8989	89FC	897E	893F	8986	8943	897D	8999	89D4	FC89	7E89	3F89	8689	4389	7D89	9989	D489	
8A	8A00	008A	8A8A	8A45	8A9A	8A4D	8ABD	8AE6	8A0B	8A73	8A02	458A	9A8A	4D8A	BD8A	E68A	0B8A	738A	028A	
8B	8B00	008B	8B8B	8BF9	8BC6	8B63	8B7B	8B85	8BF6	8BFA	8B9F	FD8B	C68B	638B	7B8B	858B	F68B	FA8B	9F8B	
8C	8C00	008C	8C8C	8C46	8C23	8CA9	8CCB	8CDD	8CE7	8CD6	8CDF	468C	238C	A98C	CB8C	DD8C	E78C	D68C	DF8C	
8D	8D00	008D	8D8D	8DFE	8D7F	8D87	8D0D	8DBE	8D1A	8D5F	8D42	FE8D	7F8D	878D	0D8D	BE8D	1A8D	5F8D	428D	
8E	8E00	008E	8E8E	8E47	8E9B	8EF5	8E36	8E1B	8E6C	8EB5	8E94	478E	9B8E	F58E	368E	1B8E	6C8E	B58E	948E	
8F	8F00	008F	8F8F	8FFF	8FC7	8FDB	8FF0	8F78	8F91	8F3C	8F09	FF8F	C78F	DB8F	F08F	788F	918F	3C8F	098F	
90	9000	0090	9090	9048	9024	9012	90E9	90CC	90A3	9066	90AE	4890	2490	1290	E990	CC90	A390	6690	AE90	
91	9100	0091	9191	91F0	9178	913C	912F	91AF	915E	91EF	9133	F091	7891	3C91	2F91	AF91	5E91	EF91	3391	
92	9200	0092	9292	9249	929C	924E	9214	920A	9228	9205	92E5	4992	9C92	4E92	1492	0A92	2892	0592	E592	
93	9300	0093	9393	93F1	93C0	9360	93D2	9369	93D5	938C	9378	F193	C093	6093	D293	6993	D593	8C93	7893	
94	9400	0094	9494	944A	9425	94AA	9462	9431	94C4	94A0	9438	4A94	2594	AA94	6294	3194	C494	A094	3894	
95	9500	0095	9595	95F2	9579	9584	95A4	9552	9539	9529	95A5	F295	7995	8495	A495	5295	3195	3995	A595	
96	9600	0096	9696	964B	969D	96F6	969F	96F7	964F	96C3	9673	4B96	9D96	F696	9F96	F796	4F96	C396	7396	
97	9700	0097	9797	97F3	97C1	97D8	9759	9794	97B2	974A	97EE	F397	C197	D897	5997	9497	B297	4A97	EE97	
98	9800	0098	9898	984C	9826	9813	988E	9847	986D	989B	98F3	4C98	2698	1398	E898	4798	6D98	9B98	F398	
99	9900	0099	9999	99F4	997A	993D	9948	9924	9990	9912	996E	F499	7A99	3D99	4899	2A99	9099	1299	6E99	
9A	9A00	009A	9A9A	9A4D	9A9E	9A4F	9A73	9A81	9AE6	9AF8	9AB8	4D9A	9E9A	4F9A	739A	819A	E69A	F89A	B89A	

Table 130: x8 ECC Correctable Syndromes

Error Bit-mask	Symbol In Error																			
	12h	11h	10h	Fh	Eh	Dh	Ch	Bh	Ah	9h	8h	7h	6h	5h	4h	3h	2h	1h	0h	
9B	9B00	009B	9B9B	9BF5	9BC2	9B61	9BB5	9BE2	9B1B	9B71	9B25	F59B	C29B	619B	B59B	E29B	1B9B	719B	259B	
9C	9C00	009C	9C9C	9C4E	9C27	9CAB	9C05	9CBA	9C0A	9C5D	9C65	4E9C	279C	AB9C	059C	BA9C	0A9C	5D9C	659C	
9D	9D00	009D	9D9D	9DF6	9D7B	9D85	9DC3	9DD9	9DF7	9DD4	9DF8	F69D	7B9D	859D	C39D	D99D	F79D	D49D	F89D	
9E	9E00	009E	9E9E	9E4F	9E9F	9EF7	9EF8	9E7C	9E81	9E3E	9E2E	4F9E	9F9E	F79E	F89E	7C9E	819E	3E9E	2E9E	
9F	9F00	009F	9F9F	9FF7	9FC3	9FD9	9F3E	9F1F	9F7C	9FB7	9FB3	F79F	C39F	D99F	3E9F	1F9F	7C9F	B79F	B39F	
A0	A000	00A0	A0A0	A050	A028	A014	A0CA	A065	A0E5	A08A	A011	50A0	28A0	14A0	CAA0	65A0	E5A0	8AA0	11A0	
A1	A100	00A1	A1A1	A1E8	A174	A13A	A10C	A106	A118	A103	A18C	E8A1	74A1	3AA1	0CA1	06A1	18A1	03A1	8CA1	
A2	A200	00A2	A2A2	A251	A290	A248	A237	A2A3	A26E	A2E9	A25A	51A2	90A2	48A2	37A2	A3A2	6EA2	E9A2	5AA2	
A3	A300	00A3	A3A3	A3E9	A3CC	A366	A3F1	A3C0	A393	A360	A3C7	E9A3	CCA3	66A3	F1A3	C0A3	93A3	60A3	C7A3	
A4	A400	00A4	A4A4	A452	A429	A44C	A441	A498	A482	A44C	A487	52A4	29A4	ACA4	41A4	98A4	82A4	4CA4	87A4	
A5	A500	00A5	A5A5	A5EA	A575	A582	A587	A5FB	A57F	A5C5	A51A	EAA5	75A5	82A5	87A5	FBA5	7FA5	C5A5	1AA5	
A6	A600	00A6	A6A6	A653	A691	A6F0	A6BC	A65E	A609	A62F	A6CC	53A6	91A6	F0A6	BCA6	5EA6	09A6	2FA6	CCA6	
A7	A700	00A7	A7A7	A7EB	A7CD	A7DE	A77A	A73D	A7F4	A7A6	A751	EBA7	CDA7	DEA7	7AA7	3DA7	F4A7	A6A7	51A7	
A8	A800	00A8	A8A8	A85A	A82A	A815	A8AD	A8EE	A8E2	A877	54A8	2AA8	15A8	ADA8	EEA8	E5A8	8AA8	77A8	4CA8	
A9	A900	00A9	A9A9	A9EC	A976	A93B	A96B	A98D	A9D6	A9FE	A9D1	ECA9	76A9	3BA9	6BA9	8DA9	D6A9	FEA9	D1A9	
AA	AA00	00AA	AAAA	AA55	AA92	AA49	AA50	AA28	AAA0	AA14	AA07	55AA	92AA	49AA	50AA	28AA	A0AA	14AA	07AA	
AB	AB00	00AB	ABAB	ABED	ABCE	AB67	AB96	AB4B	AB5D	AB9D	AB9A	EDAB	CEAB	67AB	96AB	4BAB	5DAB	9DAB	9AAB	
AC	AC00	00AC	ACAC	AC56	AC2B	ACAD	AC26	AC13	AC4C	ACB1	ACBA	56AC	2BAC	ADAC	26AC	13AC	4CAC	B1AC	DAAC	
AD	AD00	00AD	ADAD	ADEE	AD77	AD83	ADE0	AD70	ADB1	AD38	AD47	EAD0	77AD	83AD	E0AD	70AD	B1AD	38AD	47AD	
AE	AE00	00AE	AEAE	AE57	AE93	AEF1	AEDB	AED5	AEC7	AED2	AE91	57AE	93AE	F1AE	DBAE	D5AE	C7AE	D2AE	91AE	
AF	AF00	00AF	AFAF	AFEF	AFCF	AFD7	AF1D	AFB6	AF3A	AF5B	AF0C	EFAF	CFAF	DFAF	1DAF	B6AF	3AAF	5BAF	0CAF	
B0	B000	00B0	B0B0	B058	B02C	B016	B004	B002	B008	B001	B0AB	58B0	2CB0	16B0	04B0	02B0	08B0	01B0	ABB0	
B1	B100	00B1	B1B1	B1E0	B170	B138	B1C2	B161	B1F5	B188	B136	E0B1	70B1	38B1	C2B1	61B1	F5B1	88B1	36B1	
B2	B200	00B2	B2B2	B259	B294	B24A	B2F9	B2C4	B283	B262	B2E0	59B2	94B2	4AB2	F9B2	C4B2	83B2	62B2	E0B2	
B3	B300	00B3	B3B3	B3E1	B3C8	B364	B33F	B3A7	B37E	B3EB	B37D	E1B3	C8B3	64B3	3FB3	A7B3	7EB3	EBB3	7DB3	
B4	B400	00B4	B4B4	B45A	B42D	B4AE	B48F	B4FF	B46F	B4C7	B43D	5AB4	2DB4	AEB4	8FB4	A7B4	6FB4	C7B4	3DB4	
B5	B500	00B5	B5B5	B5E2	B571	B580	B549	B59C	B592	B54E	B5A0	E2B5	71B5	80B5	49B5	9CB5	92B5	4EB5	A0B5	
B6	B600	00B6	B6B6	B65B	B695	B6F2	B672	B639	B6E4	B6A4	B676	5BB6	95B6	F2B6	72B6	39B6	E4B6	A4B6	76B6	
B7	B700	00B7	B7B7	B7E3	B7C9	B7DC	B7B4	B75A	B719	B72D	B7EB	E3B7	C9B7	DCB7	B4B7	5AB7	19B7	2DB7	EBB7	
B8	B800	00B8	B8B8	B85C	B82E	B817	B863	B889	B8C6	B8FC	B8F6	5CB8	2EB8	17B8	63B8	89B8	6CB8	FCB8	6FB8	
B9	B900	00B9	B9B9	B9E4	B972	B939	B9A5	B9EA	B93B	B975	B96B	E4B9	72B9	39B9	A5B9	EAB9	3BB9	75B9	6BB9	
BA	BA00	00BA	BABA	BA5D	BA96	BA4B	BA9E	BA4F	BA4D	BA9F	BABD	5DBA	96BA	4BBA	9EBA	4FBA	4DBA	9FBA	BDBA	
BB	BB00	00BB	BBBB	BBE5	BBCA	BB65	BB58	BB2C	BBB0	BB16	BB20	E5BB	CABB	65BB	58BB	2CBB	B0BB	16BB	20BB	
BC	BC00	00BC	BCBC	BC5E	BC2F	B4AE	BCE8	BC7A	BCA1	BC3A	BC60	5EBC	2FBC	AFBC	E8BC	74BC	A1BC	3ABC	60BC	
BD	BD00	00BD	BDBD	BDE6	BD73	BD81	BD2E	BD17	BD5C	BDB3	BDFD	E6BD	73BD	81BD	2EBD	17BD	5CBD	B3BD	FDBD	
BE	BE00	00BE	BEBE	BE5F	BE97	BEF3	BE15	BEB2	BE2A	BE59	BE2B	5FBE	97BE	F3BE	15BE	B2BE	2ABE	59BE	2BBE	
BF	BF00	00BF	BFBF	BFE7	BFCB	BFDD	BFD3	BFD1	BFD7	BFD0	BFB6	E7BF	CBBF	DDBF	D3BF	D1BF	D7BF	D0BF	B6BF	
C0	C000	00C0	C0C0	C060	C030	C018	C08C	C046	C069	C023	C01E	60C0	30C0	18C0	8CC0	60C0	23C0	1EC0		
C1	C100	00C1	C1C1	C1D8	C16C	C136	C14A	C125	C194	C1AA	C183	D8C1	6CC1	36C1	4AC1	25C1	94C1	AAC1	83C1	
C2	C200	00C2	C2C2	C261	C288	C244	C271	C280	C2E2	C240	C255	61C2	88C2	44C2	71C2	80C2	E2C2	40C2	55C2	
C3	C300	00C3	C3C3	C3D9	C3D4	C36A	C3B7	C3E3	C31F	C3C9	C3C8	D9C3	D4C3	6AC3	B7C3	E3C3	1FC3	C9C3	C8C3	
C4	C400	00C4	C4C4	C462	C431	C4A0	C407	C4BB	C40E	C4E5	C488	62C4	31C4	A0C4	07C4	BBC4	0EC4	ESC4	88C4	
C5	C500	00C5	C5C5	C5DA	C56D	C58E	C5C1	C5D8	C5F3	C56C	C515	DAC5	6DC5	8EC5	C1C5	D8C5	F3C5	6CC5	15C5	
C6	C600	00C6	C6C6	C663	C689	C6FC	C6FA	C67D	C685	C686	C6C3	63C6	89C6	FCC6	FAC6	7DC6	85C6	86C6	C3C6	
C7	C700	00C7	C7C7	C7B6	C7D5	C7D2	C73C	C71E	C778	C70F	C75E	DBC7	D5C7	D2C7	3CC7	1EC7	78C7	0FC7	5EC7	
C8	C800	00C8	C8C8	C864	C832	C819	C8EB	C8CD	C8A7	C8DE	C843	62C8	32C8	19C8	EBC8	CDC8	A7C8	DEC8	43C8	
C9	C900	00C9	C9C9	C9DC	C9E6	C937	C92D	C9AE	C95A	C957	C9DE	DCC9	6EC9	37C9	2DC9	AEC9	5AC9	57C9	DEC9	
CA	CA00	00CA	CACA	CA65	CA8A	CA45	CA16	CA0B	CA2C	CABD	CA08	65CA	8ACA	45CA	16CA	0BCA	2CCA	BDCA	08CA	
CB	CB00	00CB	CBCB	CBDD	CBDE	CB6B	CBD0	CB68	CBD1	CB34	CB95	DDCB	D6CB	6BCB	D0CB	68CB	D1CB	34CB	95CB	
CC	CC00	00CC	CCCC	CC66	CC33	CCA1	CC60	CC30	CC0C	CC18	CCD5	66CC	33CC	A1CC	60CC	30CC	C0CC	18CC	55CC	
CD	CD00	00CD	CDCD	CDDE	CD6F	CD8F	CDA6	CD53	CD3D	CD91	CD48	DECD	6FCD	8FCD	A6CD	53CD	3DCD	91CD	48CD	
CE	CE00	00CE	CECE	CE67	CE8B	CEFD	CE9D	CE6F	CE4B	CE7B	CE9E	67CE	8BCE	FDCE	9DCE	F6CE	4BCE	7BCE	9ECE	
CF	CF00	00CF	CFCF	CFDF	CFD7	CFD3	CF5B	CF95	CFB6	CFE2	CF03	DFCF	D7CF	D3CF	5BCF	95CF	B6CF	F2CF	03CF	
D0	D000	00D0	D0D0	D068	D034	D01A	D042	D021	D084	D0A8	D0A4	68D0	34D0	1AD0	42D0	21D0	84D0	A8D0	A4D0	
D1	D100	00D1	D1D1	D1D0	D168	D134	D184	D142	D179	D121	D139	D0D1	68D1	34D1	84D1	42D1	79D1	21D1	39D1	
D2	D200	00D2	D2D2	D269	D28C	D246	D2BF	D2E7	D20F	D2CB	D2EF	69D2	8CD2	46D2	BFD2	E7D2	0FD2	CB2D	EFD2	
D3	D300	00D3	D3D3	D3D1	D3D0	D368	D379	D384	D3F2	D342	D372	D1D3	D0D3	68D3	79D3	84D3	F2D3	42D3	72D3	
D4	D400	00D4	D4D4	D46A	D435	D4A2	D4C9	D4DC	D4E3	D46E	D432	6AD4	35D4	A2D4	C9D4	DCD4	E3D4	6ED4	32D4	
D5	D500	00D5	D5D5	D5D2	D569	D58C	D50F	D5BF	D51E	D5E7	D5AF	D2D5	69D5	8CD5	0FD5	BFD5	1ED5	E7D5	AFD5	
D6	D600	00D6	D6D6	D66B	D68D	D6FE	D634	D61A	D668	D6D0	D679	6BD6	8DD6	FED6	34D6	1AD6	68D6	0DD6	79D6	
D7	D700	00D7	D7D7	D7D3	D7D1	D7D0	D7F2	D779	D795	D784	D7E4	D3D7	D1D7	D0D7	F2D7	79D7	95D7	84D7	E4D7	
D8	D800	00D8	D8D8	D86C	D836	D81B	D825	D8AA	D84A	D855	D8F9	6CD8	36D8	1BD8	25D8	AAD8	4AD8	55D8	F9D8	
D9	D900	00D9	D9D9	D9D4	D96A	D935	D9E3	D9C9	D9B7	D9DC	D964	D4D9	6AD9	35D9	E3D9	C9D9	B7D9	DCD9	64D9	
DA	DA00	00DA	DADA	DA6D	DA8E	DA47	DAD8	DA6C	DAC1	DA36	DAB2	6DDA	8EDA	47DA	D8DA	6CDA	C1DA	36DA	B2DA	
DB	DB00	00DB	DBDB	DBD5	DBD2	DB69	DB1E	DB0F	DB3C	DBBF	DB2F	D5DB	D2DB	69DB	1EDB	0FDB	3CDB	BFD5	2FDB	
DC	DC00	00DC	DCDC	DC6E	DC37	DCA3	DCAE	DC57	DC2D	DC93	DC6F	6EDC	37DC	A3DC	AEDC	57DC	2DDC	93DC	6FDC	
DD	DD00	00DD	DDDD	DDD6	DD6B	DD8D	DD68	DD34	DDD0	DD1A	DDF2	D6DD	6BDD	8DDD	68DD	34DD	D0DD	1ADD	F2DD	

Table 130: x8 ECC Correctable Syndromes

Error Bit-mask	Symbol In Error																			
	12h	11h	10h	Fh	Eh	Dh	Ch	Bh	Ah	9h	8h	7h	6h	5h	4h	3h	2h	1h	0h	
DE	DE00	00DE	DEDE	DE6F	DE8F	DEFF	DE53	DE91	DEA6	DEF0	DE24	6FDE	8FDE	FFDE	53DE	91DE	A6DE	F0DE	24DE	
DF	DF00	00DF	DFDF	DFD7	DFD3	DFD1	DF95	DFF2	DF5B	DF79	DFB9	D7DF	D3DF	D1DF	95DF	F2DF	5BDF	79DF	B9DF	
E0	E000	00E0	E0E0	E070	E038	E01C	E061	E088	E0C2	E044	E01B	70E0	38E0	1CE0	61E0	88E0	C2E0	44E0	1BE0	
E1	E100	00E1	E1E1	E1C8	E164	E132	E1A7	E1EB	E13F	E1CD	E186	C8E1	64E1	32E1	A7E1	EBE1	3FE1	CDE1	86E1	
E2	E200	00E2	E2E2	E271	E280	E240	E29C	E24E	E249	E227	E250	71E2	80E2	40E2	9CE2	4EE2	49E2	27E2	50E2	
E3	E300	00E3	E3E3	E3C9	E3DC	E36E	E35A	E32D	E3B4	E3AE	E3CD	C9E3	DCE3	6EE3	5AE3	A2E3	B4E3	AEE3	CDE3	
E4	E400	00E4	E4E4	E472	E439	E4A4	E4EA	E475	E4A5	E482	E48D	72E4	39E4	A4E4	EAE4	75E4	A5E4	82E4	8DE4	
E5	E500	00E5	E5E5	E5CA	E565	E58A	E52C	E516	E558	E50B	E510	CAE5	6E5E	8AE5	2CE5	16E5	58E5	0BE5	10E5	
E6	E600	00E6	E6E6	E673	E681	E6F8	E617	E6B3	E62E	E6E1	E6C6	73E6	81E6	F8E6	17E6	B3E6	2EE6	E1E6	C6E6	
E7	E700	00E7	E7E7	E7CB	E7D7	E7D6	E7D1	E7D0	E7D3	E768	E75B	CBE7	DDE7	D6E7	D1E7	D0E7	D3E7	68E7	5BE7	
E8	E800	00E8	E8E8	E874	E83A	E81D	E806	E803	E80C	E8B9	E846	74E8	3AE8	1DE8	06E8	03E8	0CE8	B9E8	46E8	
E9	E900	00E9	E9E9	E9CC	E966	E933	E9C0	E960	E9F1	E930	E9DB	CCE9	66E9	33E9	C0E9	60E9	F1E9	30E9	DBE9	
EA	EA00	00EA	EAEA	EA75	EA82	EA41	EAFB	EAC5	EA87	EADA	EA0D	75EA	82EA	41EA	FAEA	C5EA	87EA	DAEA	0DEA	
EB	EB00	00EB	EBEA	EBCD	EB6E	EB3D	EBA6	EBA5	EB53	EB90	CDEB	DEEB	6FEB	3DEB	3DEB	3DEB	3DEB	3DEB	90EB	
EC	EC00	00EC	ECEC	EC76	EC3B	ECA5	EC8D	ECFE	EC6B	EC7F	ECD0	76EC	3BEC	A5EC	8DEC	FEEC	6BEC	7FEC	D0EC	
ED	ED00	00ED	EDED	EDCE	ED67	ED8B	ED4B	ED9D	ED96	EDF6	ED4D	CEED	67ED	8BED	4BED	9DED	96ED	F6ED	4DED	
EE	EE00	00EE	EEEE	EE77	EE83	EEF9	EE70	EE38	EEEE	EE1C	EE9B	77EE	83EE	F9EE	70EE	38EE	E0EE	1CEE	9BEE	
EF	EF00	00EF	EFEF	EFCF	EFD7	EFD7	EFB6	EF5B	EF1D	EF95	EF06	CFEF	DFEF	D7EF	B6EF	5BEF	1DEF	95EF	06EF	
F0	F000	00F0	F0F0	F078	F03C	F01E	F0AF	F0EF	F02F	F0CF	F0A1	78F0	3CF0	1EF0	AF00	EFF0	2FF0	CFF0	A1F0	
F1	F100	00F1	F1F1	F1C0	F160	F130	F169	F18C	F1D2	F146	F13C	C0F1	60F1	30F1	69F1	8CF1	D2F1	46F1	3CF1	
F2	F200	00F2	F2F2	F279	F284	F242	F252	F229	F2A4	F2AC	F2EA	79F2	84F2	42F2	52F2	29F2	A4F2	ACF2	EAF2	
F3	F300	00F3	F3F3	F3C1	F3D8	F36C	F394	F34A	F359	F325	F377	C1F3	D8F3	6CF3	94F3	4AF3	59F3	25F3	77F3	
F4	F400	00F4	F4F4	F47A	F43D	F4A6	F424	F412	F448	F409	F437	7AF4	3DF4	A6F4	24F4	12F4	48F4	09F4	37F4	
F5	F500	00F5	F5F5	F5C2	F561	F588	F5E2	F571	F5B5	F580	F5AA	C2F5	61F5	88F5	E2F5	71F5	B5F5	80F5	AAF5	
F6	F600	00F6	F6F6	F67B	F685	F6FA	F6D9	F6D4	F6C3	F66A	F67C	7BF6	85F6	FAF6	D9F6	D4F6	C3F6	6AF6	7CF6	
F7	F700	00F7	F7F7	F7C3	F7D9	F7D4	F71F	F7B7	F73E	F7E3	F7E1	C3F7	D9F7	D4F7	1FF7	B7F7	3EF7	E3F7	E1F7	
F8	F800	00F8	F8F8	F87C	F83E	F81F	F8C8	F864	F8E1	F832	F8FC	7CF8	3EF8	1FF8	C8F8	64F8	E1F8	32F8	FCF8	
F9	F900	00F9	F9F9	F9C4	F962	F931	F90E	F907	F91C	F9BB	F961	C4F9	62F9	31F9	0EF9	07F9	1CF9	BBF9	61F9	
FA	FA00	00FA	FAFA	FA7D	FA86	FA43	FA35	FAA2	FA6A	FA51	FAB7	7DFA	86FA	43FA	35FA	A2FA	6AFA	51FA	B7FA	
FB	FB00	00FB	FBFB	FBC5	FBDA	FB6D	FBF3	FBC1	FB97	FB88	FB2A	C5FB	DAFB	6DFB	F3FB	C1FB	97FB	D8FB	2AFB	
FC	FC00	00FC	FCFC	FC7E	FC3F	FCA7	FC43	FC99	FC86	FCF4	FC6A	7EFC	3FFC	A7FC	43FC	99FC	86FC	F4FC	6AFC	
FD	FD00	00FD	FD7D	FDC6	FD63	FD89	FD85	FDFA	FD7B	FD7D	FD77	C6FD	63FD	89FD	85FD	FAFD	7BFD	7DFD	F7FD	
FE	FE00	00FE	FEFE	FE7F	FE87	FEFB	FEBE	FE5F	FE0D	FE97	FE21	7FFE	87FE	FBFE	BEFE	5FFE	0DFE	97FE	21FE	
FF	FF00	00FF	FFFF	FFC7	FFDB	FFD5	FF78	FF3C	FFF0	FF1E	FFBC	C7FF	DBFF	D5FF	78FF	3CFF	F0FF	1EFF	BCFF	

2.13.2.5.2 x4 ECC

The use of x4 ECC is indicated in [D18F3x180](#)[EccSymbolSize].

The syndrome field uniquely identifies the failing bit positions of a correctable ECC error. Only syndromes identified by [Table 131](#) are correctable by the error correcting code.

Symbols 00h-0Fh map to data bits 0-63; symbols 10h-1Fh map to data bits 64-127; symbols 20-21h map to ECC check bits for data bits 0-63; symbols 22-23h map to ECC check bits for data bits 64-127.

To use [Table 131](#), first find the 16-bit syndrome value in the table. This is performed by using low order 4 bits of the syndrome to select the appropriate error bitmask column. The entire four digit syndrome should then be in one of the rows of that column. The Symbol In Error row indicates which symbol, and therefore which DIMM has the error, and the column indicates which bits within the symbol.

For example, if the ECC syndrome is 6913h, then symbol 05h has the error, and bits 0 and 1 within that symbol are corrupted, since the syndrome is in column 3h (0011b). Symbol 05h maps to bits 23-20, so the corrupted bits are 20 and 21.

Table 131: x4 ECC Correctable Syndromes

Symbol In Error	Error Bitmask														
	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011	1100	1101	1110	1111
Data 0	e821	7c32	9413	bb44	5365	c776	2f57	dd88	35a9	a1ba	499b	66cc	8eed	1afe	f2df
Data 1	5d31	a612	fb23	9584	c8b5	3396	6ea7	eac8	b7f9	4cda	11eb	7f4c	227d	d95e	846f
Data 2	0001	0002	0003	0004	0005	0006	0007	0008	0009	000a	000b	000c	000d	000e	000f
Data 3	2021	3032	1013	4044	6065	7076	5057	8088	a0a9	b0ba	909b	c0cc	e0ed	f0fe	d0df
Data 4	5041	a082	f0c3	9054	c015	30d6	6097	e0a8	b0e9	402a	106b	70fc	20bd	d07e	803f
Data 5	be21	d732	6913	2144	9f65	f676	4857	3288	8ca9	e5ba	5b9b	13cc	aded	c4fe	7adf
Data 6	4951	8ea2	c7f3	5394	1ac5	dd36	9467	a1e8	e8b9	2f4a	661b	f27c	bb2d	7cde	358f
Data 7	74e1	9872	ec93	d6b4	a255	4ec6	3a27	6bd8	1f39	f3aa	874b	bd6c	c98d	251e	51ff
Data 8	15c1	2a42	3f83	cef4	db35	e4b6	f177	4758	5299	6d1a	78db	89ac	9c6d	a3ee	b62f
Data 9	3d01	1602	2b03	8504	b805	9306	ae07	ca08	f709	dc0a	e10b	4f0c	720d	590e	640f
Data 10	9801	ec02	7403	6b04	f305	8706	1f07	bd08	2509	510a	c90b	d60c	4e0d	3a0e	a20f
Data 11	d131	6212	b323	3884	e9b5	5a96	8ba7	1cc8	cdf9	7eda	afeb	244c	f57d	465e	976f
Data 12	e1d1	7262	93b3	b834	59e5	ca56	2b87	dc18	3dc9	ae7a	4fab	642c	85fd	164e	f79f
Data 13	6051	b0a2	d0f3	1094	70c5	a036	c067	20e8	40b9	904a	f01b	307c	502d	80de	e08f
Data 14	a4c1	f842	5c83	e6f4	4235	1eb6	ba77	7b58	df99	831a	27db	9dac	396d	65ee	c12f
Data 15	11c1	2242	3383	c8f4	d935	eab6	fb77	4c58	5d99	6e1a	7fdb	84ac	956d	a6ee	b72f
Data 16	45d1	8a62	cfb3	5e34	1be5	d456	9187	a718	e2c9	2d7a	68ab	f92c	bcfd	734e	369f
Data 17	63e1	b172	d293	14b4	7755	a5c6	c627	28d8	4b39	99aa	fa4b	3c6c	5f8d	8d1e	eeff
Data 18	b741	d982	6ec3	2254	9515	fbd6	4c97	33a8	84e9	ea2a	5d6b	11fc	a6bd	c87e	7f3f
Data 19	dd41	6682	bbc3	3554	e815	53d6	8e97	1aa8	c7e9	7c2a	a16b	2ffc	f2bd	497e	943f
Data 20	2bd1	3d62	16b3	4f34	64e5	7256	5987	8518	aec9	b87a	93ab	ca2c	e1fd	f74e	dc9f
Data 21	83c1	c142	4283	a4f4	2735	65b6	e677	f858	7b99	391a	badb	5cac	df6d	9dee	1e2f
Data 22	8fd1	c562	4ab3	a934	26e5	6c56	e387	fe18	71c9	3b7a	b4ab	572c	d8fd	924e	1d9f
Data 23	4791	89e2	ce73	5264	15f5	db86	9c17	a3b8	e429	2a5a	6dcb	f1dc	b64d	783e	3faf
Data 24	5781	a9c2	fe43	92a4	c525	3b66	6ce7	e3f8	b479	4a3a	1dbb	715c	26dd	d89e	8f1f
Data 25	bf41	d582	6ac3	2954	9615	fcd6	4397	3ea8	81e9	eb2a	546b	17fc	a8bd	c27e	7d3f
Data 26	9391	e1e2	7273	6464	f7f5	8586	1617	b8b8	2b29	595a	cacb	dcdc	4f4d	3d3e	aeaf
Data 27	cce1	4472	8893	fdb4	3155	b9c6	7527	56d8	9a39	12aa	de4b	ab6c	678d	ef1e	23ff
Data 28	a761	f9b2	5ed3	e214	4575	1ba6	bcc7	7328	d449	8a9a	2dfb	913c	365d	688e	cfef
Data 29	ff61	55b2	aad3	7914	8675	2ca6	d3c7	9e28	6149	cb9a	34fb	e73c	185d	b28e	4def
Data 30	5451	a8a2	fcf3	9694	c2c5	3e36	6a67	ebe8	bf9b	434a	171b	7d7c	292d	d5de	818f
Data 31	6fc1	b542	da83	19f4	7635	acb6	c377	2e58	4199	9b1a	f4db	37ac	586d	82ee	ed2f
Check0	be01	d702	6903	2104	9f05	f606	4807	3208	8c09	e50a	5b0b	130c	ad0d	c40e	7a0f
Check1	4101	8202	c303	5804	1905	da06	9b07	ac08	ed09	2e0a	6f0b	f40c	b50d	760e	370f
Check2	c441	4882	8cc3	f654	3215	bed6	7a97	5ba8	9fe9	132a	d76b	adfc	69bd	e57e	213f

Table 131: x4 ECC Correctable Syndromes

Symbol In Error	Error Bitmask														
	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011	1100	1101	1110	1111
Check3	7621	9b32	ed13	da44	ac65	4176	3757	6f88	19a9	f4ba	829b	b5cc	c3ed	2efe	58df

2.13.3 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:

- DRAM: [2.13.3.1 \[DRAM Error Injection\]](#).
- Links:
 - [D18F0x\[14C:130\]\[ForceRetryError\]](#), [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.

2.13.3.1 DRAM Error Injection

This section gives details and examples on injecting errors into DRAM using [D18F3xBC_x8 \[DRAM ECC\]](#). The intent of DRAM error injection is to cause a discrepancy between the stored data and the stored ECC value. Therefore, DRAM error injection is only possible on DRAM which supports ECC, and in which [D18F2x90_dct\[1:0\]\[DimmEccEn\]](#) and [D18F3x44\[DramEccEn\]](#) are set.

The memory subsystem operates on 64-byte cachelines. The following fields are used to set how the cacheline is to be corrupted in DRAM:

- [D18F3xB8\[ArrayAddress\]](#) selects a cacheline quadrant (16-byte section) of the cacheline. Each cacheline quadrant is protected by an ECC word. Note that there are special requirements for which bits are used to specify the target quadrant.
- [D18F3xBC_x8\[ErrInjEn\]](#) selects a 16-bit word of the cacheline quadrant selected in ArrayAddress. The 16-bit word identified as ECC[15:0] refers to the bits which store the ECC value; the other 16-bit words address the data on which the ECC is calculated. One or more of these 16-bit words can be selected, and the error bitmask indicated in ErrVector is applied to each of the selected words.
- [D18F3xBC_x8\[ErrVector\]](#) is a bitmask which selects the individual bits to be corrupted in the 16-bit words selected by ErrInjEn. When selecting the bits to be corrupted for correctable or uncorrectable errors, consider the ECC scheme being used, including symbol size; see [2.13.2 \[DRAM Considerations for ECC\]](#) for more details. Note that corrupting more than two symbols may exceed the guarantees of the ECC to detect the errors; for testing purposes it is recommended that no more than two symbols be corrupted in a single cacheline quadrant.

The distinction between [D18F3xBC_x8\[DramErrEn\]](#) and [D18F3xBC_x8\[EccWrReq\]](#) is that DramErrEn is used to continuously inject errors on every write. This bit is set and cleared by software. EccWrReq is used to inject an error on only one write. This bit is set by software and is cleared by hardware after the error is injected.

When performing DRAM error injection on multi-node systems, the error injection registers of the NB to which the memory is attached must be programmed.

The following can be used to trigger the injection:

- The memory address is not an explicit parameter of the error injection interface. Once the error injection registers [D18F3xB8](#) and [D18F3xBC](#) are set, the next non-cached access of the appropriate type will trigger the mechanism and apply it to the accessed address. The access should be non-cached so that it is guaranteed to be seen by the memory controller. Possible methods to ensure a non-cached access include using the appropriate MTRR to set the memory type to UC or turning off caches. If it is important to know the address, then system activity must be quiesced so that the access can take place under careful software control. Once the error injection pattern is set in [D18F3xB8](#) and [D18F3xBC_x8](#):
 - Set either [D18F3xBC_x8\[EccWrReq\]](#) or [D18F3xBC_x8\[DramErrEn\]](#) to enable the triggering mechanism.
 - The next non-cached access of the appropriate type will trigger the mechanism and apply it to the accessed address.
 -

After the error is injected, the data must be referenced in order for the error detection to be triggered. The error address logged in [MSR0000_0412 \[NB Machine Check Address \(MC4_ADDR\)\]](#) will correspond to the cache-line quadrant that contains the error.

When using [MSR0000_0411](#) to read [MC4_STATUS](#) after an error injection and subsequent error detection, be aware that the setting of [D18F3x44\[NbMcaToMstCpu\]](#) can cause different cores to see different values. Alternatively, [MC4_STATUS](#) can be read through the PCI-defined configuration space aliases [D18F3x4C](#) and [D18F3x48](#), which do not return different values to different cores, regardless of the setting of [D18F3x44\[NbMcaToMstCpu\]](#).

Example 1: Injecting a correctable error:

- Program error pattern:
 - [D18F3xB8\[ArraySelect\]](#)=1000b // select DRAM as target
 - [D18F3xB8\[ArrayAddress\]](#)=000000000b // select 16-byte (128-bit) section
 - [D18F3xBC_x8\[ErrInjEn\]](#)=000000001b // select 16-bit word in 16-byte section
 - [D18F3xBC_x8\[EccRdReq\]](#)=0 // not a read request
 - [D18F3xBC_x8\[EccVector\]](#)=0001h // set bitmask to inject error into only one symbol
- Program error trigger:
 - [D18F3xBC_x8\[DramErrEn\]](#)=0 // inject only a single error
 - [D18F3xBC_x8\[EccWrReq\]](#)=1 // a write request; enable injection on next write
- Clean up // if programmed for continuous errors
 - [D18F3xBC_x8\[DramErrEn\]](#)=0 // inject only a single error

2.13.4 Sideband Interface (SBI)

The sideband interface (SBI) is an SMBus v2.0 compatible 2-wire processor slave interface. SBI is also referred as the Advanced Platform Management Link. All I2C v2.1 speeds are supported.

SBI is used to communicate with the Remote Management Interface (SB-RMI) (see the *Advanced Platform Management Link (APML) Specification*, #41918). The processor is compliant to APML Revision=03h ([D18F3x1EC_x100\[Revision\]](#)==03h).

2.13.4.1 SBI Processor Information

Processor access to the SBI configuration is via [D18F3x1E4 \[SBI Control\]](#). The processor can access the SB-RMI registers through [D18F3x1E8 \[SBI Address\]](#) and [D18F3x1EC \[SBI Data\]](#).

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 **IOCFC [IO-Space Configuration Data Port]**. Accesses to these registers from each core of a processor target the same registers of that processor; it is not possible for a node to access these registers on a different node.
- **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[_xYYY]**: processor capabilities information returned by the CPUID instruction. See 3.10 [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. E.g., D18F3x40 specifies the register at device 18h, function 3, and address 40h. See 2.8 [Configuration Space], for details about configuration space.
 - X specifies the hexadecimal device number; this may be 1 or 2 digits.
 - Registers specified as D18:
 - Unless otherwise specified there is one set of these registers **Per-node**; these registers in any node are accessible through any core of any node.
 - Bus 0 Device24 is D18=Node0; D18F0x60[NodeID]=0.
 - Bus 0 Device25 is D19=Node1; D19F0x60[NodeID]=1.
 - ...
 - Bus 0 Device31 is D1F=Node7; D1Fx60[NodeID]=7.
 - Y specifies the function number; this is 1 digit ranging in value from 0-7.
 - ZZZ specifies the hexadecimal byte address; this may be 2 or 3 digits.
 - Some registers in D18F2xXXX have the **_dct[1:0]** mnemonic suffix. See 2.10.1 [DCT Configuration Registers].
 - Some registers accessed via D18F4x1[98,90,88,80] [Link Phy Offset] have the “**_dm[1]**” mnemonic suffix. See D18F4x1[98,90,88,80] [Link Phy Offset].
- **PMCxXXX**: core 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 2.7.1 [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 2.7.2 [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 spec-

ified 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 a 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 132: Terminology in Register Descriptions

Term	Definition
BIOS	Software recommendation syntax. See 3.1.2 [Software Recommendation (BIOS, SBIOS)].
SBIOS	
See	Reference to remote definition. See 3.1.3 [See Keyword (See:)].
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.

Table 132: Terminology in Register Descriptions

Term	Definition
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	
SBI-only	No host software access; host accesses result in GP-read-write. HDT access supported. SB-RMI access documented by the APML spec. See 2.13.4 [Sideband Interface (SBI)].
Reset-applied	Takes affect on warm reset.
GP-read	GP exception occurs on read.
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. 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: SameOnAllCheckEnabled = ((same-for-all MSR CPUID) && ~(not-same-for-all UpdatedByHw)). UpdatedByHw = (Updated-by-hardware set-by-hardware cleared-by-hardware set-when-done cleared-when-done).
Same-for-all	
Same-within-node	
Same-within-pkg	
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	

Table 132: Terminology in Register Descriptions

Term	Definition
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. [MSRC001_001F](#)). 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)

The following keywords specify the recommended value to be set by software.

- BIOS is AMD BIOS.
- SBIOS is system BIOS.

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 See Keyword (See:)

There is a special meaning applied to the use of “See:” that differs from the use of See not followed by a “:”.

- See, not followed by a “:”, simply refers the reader to a document location that contains related information.
- See followed by a “:” is a shorthand notation that indicates that the definition for this register or register field inherits all properties and definitions from the register or register field that follows "See:". Any definition local to the register or register field supercedes this inheritance.

“See:” can be used in the following ways:

- Full register width. [CPUID Fn0000_0001_EAX](#) inherits it's full register width definition from [D18F3xFC](#).
- Register field. [MSR0000_0277\[PA1MemType\]](#) inherits it's definition from PA0MemType, however, the local reset of 4h overrides the inherited PA0MemType reset of 6h.
- Valid values definition. [MSR0000_020\[F:0\]\[MemType\]](#), for example, inherits the valid values definition from [Table 265 \[Valid Values for Memory Type\]](#).

3.1.4 Mapping Tables

The following mapping table types are defined.

3.1.4.1 Register Mapping

The register mapping table specifies the specific function for each register in a range of registers.

[Table 133](#), for example, specifies that the D18F0x40 function is for Node 0, resulting in the D18F0x40 name “Routing Table - Node 0”.

3.1.4.2 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 191](#), for example, specifies that the D18F2x98_dct[1:0][31:0]==0D0F_0031h, or D18F2x9C_x0D0F_0031_dct[1:0], function is for Byte 0, resulting in the register name “Data Byte Fence2 Threshold - Byte0”.

3.1.4.3 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 rows. The cell at the intersection of the register row and the field bit range column specifies the suffix that is appended to the register field with a “_”. “Reserved” specifies that the field is reserved for the register of that row.

[Table 167](#), for example, specifies that the fields at D18F2x9C_x0000_0[3:0]01[31:24] should have the suffix Byte3, resulting in WrDatGrossDly_Byte3 and WrDatFineDly_Byte3.

3.1.4.4 Broadcast Mapping

The broadcast mapping table maps a register address to a range of register addresses that are written as a group when the broadcast register address is written. 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 written as a group when the row and column address is written.

[Table 192](#), for example, specifies that a write to D18F2x98_dct[1:0][31:0]==0D0F_0F31h will result in a broadcast write to the D18F2x9C_x0D0F_0[8:0]31 range of registers.

3.1.4.5 Reset Mapping

The reset mapping table specifies the reset, cold reset, or value for each register in a range of registers.

[Table 260](#), for example, specifies that the CPUID Fn0000_0000_EBX register has a value of 6874_7541h, with a comment of “The ASCII characters “h t u A””.

3.1.4.6 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 165 \[Valid Values for D18F2x94_dct\[1:0\]\[MemClkFreq\]\]](#))

3.1.4.7 BIOS Recommendations

The BIOS recommendations table defines “BIOS:” recommendations that are conditional and complex enough to warrant a table.

[Table 220 \[BIOS Recommendations for D18F3x140\[FreeTok, UpRspTok\]\]](#), for example, specifies the BIOS recommendations for [D18F3x140\[FreeTok, UpRspTok\]](#). All cells under the “Condition” header for a given row are ANDed to form the condition for the values to the right of the condition. For example, rows 1-3 provides the following equivalent BIOS recommendations:

- [D18F3x140\[FreeTok\]](#): BIOS: IF (SCM) THEN 10 ELSEIF ((MCM1 | MCM2h) & PrbFltrEn) THEN 9 ELSEIF ((MCM1 | MCM2h) & ~PrbFltrEn) THEN 10 ESEIF etc.
- [D18F3x140\[UpRspTok\]](#): BIOS: IF (SCM) THEN 3 ELSEIF ((MCM1 | MCM2h) & PrbFltrEn) THEN 3 ELSEIF ((MCM1 | MCM2h) & ~PrbFltrEn) THEN 3 ESEIF etc.

[Table 93 \[BIOS Recommendations for F0RC13\[NumLogicalRanks\]\]](#), for example, describes a condition in a different manner. When the condition column has a name, such as [NumDimmSlots](#), then only one of the list of values must match to be true. E.g. If the cell for the [NumDimmSlots](#) column contains “1, 2” then the resulting condition is (([NumDimmSlots](#)==1) | ([NumDimmSlots](#)==2)). Also, a “-” indicates that the cell is not included into the condition. The resulting BIOS recommendation for row 3 is:

- [F0RC13\[NumLogicalRanks\]](#): BIOS: ... ELSEIF (([F0RC13\[NumPhysicalRanks\]](#)==10b) & (([NumDimmSlots](#)==1) | ([NumDimmSlots](#)==2))) THEN 10b ELSEIF ...

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 [IOCF8 \[IO-Space Configuration Data Port\]](#), are used to access system configuration space, as defined by the PCI specification. [IOCF8](#) provides the address register and [IOCF8](#) provides the data port. Software sets up the configuration address by writing to [IOCF8](#). Then, when an access is made to [IOCF8](#), the processor generates the corresponding configuration access to the address specified in [IOCF8](#). See [2.8 \[Configuration Space\]](#).

[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 [IOCF8](#) 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 [IOCF8](#) 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 IOCF8 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.

IOCF8 IO-Space Configuration Data Port

Bits	Description
31:0	Data. Read-write. Reset: 0. See IOCF8 .

3.3 Device [1F:18]h Function 0 Configuration Registers

See [3.1 \[Register Descriptions and Mnemonics\]](#). See [2.8 \[Configuration Space\]](#).

D18F0x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1600h.
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: Product-specific. Specifies the offset of the link capabilities block based on the links that are supported by the node.																				
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7Fh-00h</td> <td>Reserved</td> </tr> <tr> <td>80h</td> <td>If link 0 is supported.</td> </tr> <tr> <td>9Fh-81h</td> <td>Reserved</td> </tr> <tr> <td>A0h</td> <td>If link 0 is not supported and link 1 is supported.</td> </tr> <tr> <td>BFh-A1h</td> <td>Reserved</td> </tr> <tr> <td>C0h</td> <td>If link 0 and 1 are not supported and link 2 is supported.</td> </tr> <tr> <td>DFh-C1h</td> <td>Reserved</td> </tr> <tr> <td>E0h</td> <td>If link 0, 1, and 2 are not supported and link 3 is supported.</td> </tr> <tr> <td>FFh-E1h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	7Fh-00h	Reserved	80h	If link 0 is supported.	9Fh-81h	Reserved	A0h	If link 0 is not supported and link 1 is supported.	BFh-A1h	Reserved	C0h	If link 0 and 1 are not supported and link 2 is supported.	DFh-C1h	Reserved	E0h	If link 0, 1, and 2 are not supported and link 3 is supported.	FFh-E1h	Reserved
Bits	Description																				
7Fh-00h	Reserved																				
80h	If link 0 is supported.																				
9Fh-81h	Reserved																				
A0h	If link 0 is not supported and link 1 is supported.																				
BFh-A1h	Reserved																				
C0h	If link 0 and 1 are not supported and link 2 is supported.																				
DFh-C1h	Reserved																				
E0h	If link 0, 1, and 2 are not supported and link 3 is supported.																				
FFh-E1h	Reserved																				

D18F0x[5C:40] Routing Table

Reset: 0004_0201h. As each packet is processed by the node, it is routed to the appropriate links, or remains in the node that is processing the packet, based on the source/destination node and the type of packet being processed. The destination of requests and responses determines which of these eight registers is used to route the packet; the source of probes and broadcasts determines which of these eight registers is used to route the packet. Once the routing table register is identified, the packet is routed to the destinations based on the state of the field (in that routing table register) that corresponds to the packet type.

Table 133: Register Mapping for D18F0x[5C:40]

Register	Function
D18F0x40	Node 0
D18F0x44	Node 1
D18F0x48	Node 2
D18F0x4C	Node 3
D18F0x50	Node 4
D18F0x54	Node 5
D18F0x58	Node 6
D18F0x5C	Node 7

Bits	Description
31:27	Reserved.
26:18	BCRoute: broadcast route. Specifies the routing information for broadcasts and probes. See: RQRoute.

17:9	RPRRoute: response route. Specifies the routing information for responses. See: RQRRoute.																				
8:0	RQRRoute: request route. Read-write. Specifies the routing information for requests.																				
	<table border="1"> <thead> <tr> <th><u>Bit</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Route to this node.</td> </tr> <tr> <td>[1]</td> <td>Route to link 0, Sublink 0.</td> </tr> <tr> <td>[2]</td> <td>Route to link 1, Sublink 0.</td> </tr> <tr> <td>[3]</td> <td>Route to link 2, Sublink 0.</td> </tr> <tr> <td>[4]</td> <td>Route to link 3, Sublink 0.</td> </tr> <tr> <td>[5]</td> <td>Route to link 0, Sublink 1.</td> </tr> <tr> <td>[6]</td> <td>Route to link 1, Sublink 1.</td> </tr> <tr> <td>[7]</td> <td>Route to link 2, Sublink 1.</td> </tr> <tr> <td>[8]</td> <td>Route to link 3, Sublink 1.</td> </tr> </tbody> </table>	<u>Bit</u>	<u>Description</u>	[0]	Route to this node.	[1]	Route to link 0, Sublink 0.	[2]	Route to link 1, Sublink 0.	[3]	Route to link 2, Sublink 0.	[4]	Route to link 3, Sublink 0.	[5]	Route to link 0, Sublink 1.	[6]	Route to link 1, Sublink 1.	[7]	Route to link 2, Sublink 1.	[8]	Route to link 3, Sublink 1.
<u>Bit</u>	<u>Description</u>																				
[0]	Route to this node.																				
[1]	Route to link 0, Sublink 0.																				
[2]	Route to link 1, Sublink 0.																				
[3]	Route to link 2, Sublink 0.																				
[4]	Route to link 3, Sublink 0.																				
[5]	Route to link 0, Sublink 1.																				
[6]	Route to link 1, Sublink 1.																				
[7]	Route to link 2, Sublink 1.																				
[8]	Route to link 3, Sublink 1.																				

D18F0x60 Node ID

Bits	Description																				
31:21	Reserved.																				
20:16	CpuCnt[4:0]: CPU count bits[4:0]. Read-write. CpuCnt[7:0] = {D18F0x160[CpuCnt[7:5]], D18F0x60[CpuCnt[4:0]]}. Reset: 0. CpuCnt[7:0] specifies the number of cores to be enabled in the system (the boot core of all nodes plus those cores enabled through D18F0x1DC[CpuEn]). 00h=1 core...3F = 64 cores; 40h through FFh are reserved. This field matches D18F0x60[NodeCnt] if each nodes in the system has one core; otherwise, it would be greater than D18F0x60[NodeCnt].																				
15	Reserved.																				
14:12	LkNode: lock node ID bits. Read-write. Reset: 0. Specifies the node ID of the node that contains the lock controller.																				
	<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>Node 0</td> <td>100b</td> <td>Node 4</td> </tr> <tr> <td>001b</td> <td>Node 1</td> <td>101b</td> <td>Node 5</td> </tr> <tr> <td>010b</td> <td>Node 2</td> <td>110b</td> <td>Node 6</td> </tr> <tr> <td>011b</td> <td>Node 3</td> <td>111b</td> <td>Node 7</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>	000b	Node 0	100b	Node 4	001b	Node 1	101b	Node 5	010b	Node 2	110b	Node 6	011b	Node 3	111b	Node 7
<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>																		
000b	Node 0	100b	Node 4																		
001b	Node 1	101b	Node 5																		
010b	Node 2	110b	Node 6																		
011b	Node 3	111b	Node 7																		
11	Reserved.																				
10:8	SbNode: Southbridge (IO hub) node ID bits. Read-write. Reset: 0. Specifies the node ID of the node that owns the link that connects to the system IO hub.																				
	<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>Node 0</td> <td>100b</td> <td>Node 4</td> </tr> <tr> <td>001b</td> <td>Node 1</td> <td>101b</td> <td>Node 5</td> </tr> <tr> <td>010b</td> <td>Node 2</td> <td>110b</td> <td>Node 6</td> </tr> <tr> <td>011b</td> <td>Node 3</td> <td>111b</td> <td>Node 7</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>	000b	Node 0	100b	Node 4	001b	Node 1	101b	Node 5	010b	Node 2	110b	Node 6	011b	Node 3	111b	Node 7
<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>																		
000b	Node 0	100b	Node 4																		
001b	Node 1	101b	Node 5																		
010b	Node 2	110b	Node 6																		
011b	Node 3	111b	Node 7																		
7	Reserved.																				

6:4	NodeCnt: node count bits. Read-write. Reset: 0. Specifies the number of coherent nodes in the system. Hardware only allows values to be programmed into this field that are consistent with the multi-processor capabilities of the device, as specified in D18F3xE8 [Northbridge Capabilities][MpCap] . Attempts to write values inconsistent with the capabilities of the processor result in this field not being updated.			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	000b	1 Node	100b	5 Nodes
	001b	2 Nodes	101b	6 Nodes
	010b	3 Nodes	110b	7 Nodes
	011b	4 Nodes	111b	8 Nodes
3	Reserved.			
2:0	NodeId: node ID bits. Read-write. Specifies the node ID of the node. Reset: It is reset to 0h for the boot strap processor (BSP); it is reset to 7h for all other nodes. It is expected that system configuration software programs the Node ID. The node IDs must be contiguous. E.g. {0, 1, 2, 3}; not {0, 1, 3, 4}. See APIC20[ApicId] .			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	000b	Node 0	100b	Node 4
	001b	Node 1	101b	Node 5
	010b	Node 2	110b	Node 6
	011b	Node 3	111b	Node 7

D18F0x64 Unit ID

Reset: 0000_00E0h.

Bits	Description
31:11	Reserved.
10:8	SbLink: Southbridge (IO hub) link ID. Read-write; set-by-hardware. Specifies the link to which the system IO hub is connected. It is only used by the node which owns the IO hub, as indicated in D18F0x60[NodeId] . For SbLink[1:0]: 00b=link 0; 01b=link 1; 10b=link 2; 11b=link 3. If the link is unganged, then SbLink[2] specifies the sublink: 0=sublink 0; 1=sublink 1. If the link is ganged, SbLink[2] is required to be low.
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	EnPReqHiPriTblWlk: isoc table walk enable for posted requests. Read-write. Reset: 0. 1=Enables the use of the Isoc channel for GART table walk requests issued for base channel posted requests. To use the Isoc channel for GART requests ICFM must be enabled, one D18F3x1[54,50,4C,48][IsocReqTok] must be allocated on each link that can receive GART table walk requests, and one D18F3x1[54,50,4C,48][IsocRspTok] must be allocated on each link that can receive GART table walk responses.

30:26	Reserved.										
25	CHtExtAddrEn: coherent link extended address enable. IF (D18F3xE8[MpCap]) THEN Read-write; Same-for-all. ELSE Read-only; Same-for-all. ENDIF. Reset: 0. 1=The coherent fabric supports physical addresses of greater than 40 bits. 0=Requests to addresses above 1 TB result in a master abort.										
24	DispRefModeEn. Read-write. Reset: 0. 1=Enables support for display-refresh ordering rules. BIOS must not set this bit until display-refresh buffers have been allocated and a warm reset has occurred. See 2.9.3.2.5 [Display Refresh And IFCM].										
23	InstallStateS. Read-write. Reset: 0. 1=Forces the default read block (RdBk) install state to be shared instead of exclusive.										
22:21	DsNpReqLmt: downstream non-posted request limit. Read-write. Reset: 00b. BIOS: 10b. 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><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	00b	No limit	01b	limited to 1	10b	limited to 4	11b	limited to 8
<u>Bits</u>	<u>Description</u>										
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. 1=Configuration accesses that (1) normally map to the configuration space within another node in the coherent fabric and (2) target a non-existent node as specified by D18F0x60[NodeCnt] are sent to an IO link instead. This bit should be set by BIOS once coherent fabric initialization is complete. Failure to do so may result in PCI configuration accesses to nonexistent nodes being sent into the coherent fabric, causing the system to hang.										

14:13	<p>BufRelPri: buffer release priority select. Read-write. Reset: 00b. BIOS: 01b. Specifies the number of link DWs sent while a buffer release is pending before the buffer release is inserted into the command/data stream of a busy link.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>64 DWs</td> </tr> <tr> <td>01b</td> <td>16 DWs</td> </tr> <tr> <td>10b</td> <td>8 DWs</td> </tr> <tr> <td>11b</td> <td>2 DWs</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	64 DWs	01b	16 DWs	10b	8 DWs	11b	2 DWs
<u>Bits</u>	<u>Description</u>										
00b	64 DWs										
01b	16 DWs										
10b	8 DWs										
11b	2 DWs										
12	<p>ATMMoDeEn: accelerated transition to modified mode enable. Read-write; Same-for-all. Reset: 0. BIOS: See 2.9.4.2. 1= Enable Accelerated Transition to Modified protocol. This mode enables usage of a new state MuW (ModifiedUnWritten) to the existing MOESI protocol. Must be programmed to the same state as D18F3x1B8[L3ATMMoDeEn]. Recommend disabled if (Dual-CoreEnabled==0).</p>										
11	<p>RespPassPW: response PassPW. Read-write. Reset: 0. BIOS: 1. 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.</p>										
10	<p>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 (uniprocessor & single core).</p>										
9	<p>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.</p>										
8	<p>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.</p>										
7	<p>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.</p>										
6	<p>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.</p>										
5	Reserved.										
4	<p>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.</p>										
3:0	Reserved.										

D18F0x6C Link Initialization Control

Bits	Description														
31:24	Reserved.														
23	TxSSBusPwrSaveEn: transmit source synchronous bus power save enable. Read-write; reset-applied. Cold reset: 0. BIOS: 1. 1=Don't toggle the transmit source synchronous clock when transmit data isn't valid. 0=Always toggle the transmit source synchronous clock. After this field is written to by software, the change is not applied until either a warm reset or a link disconnect sequence occurs through LDTSTOP_L.														
22:21	Reserved.														
20	Tr0Double: training 0 time double. Read-write. Cold reset: 0. 1=All Training 0 times are doubled from their programmed or default/reset values. This bit should be set when using link BIST on this device connected to an ILM device, and left clear at all other times.														
19:12	Reserved.														
11	DefLnk[2]: default link. See: DefLnk[1:0].														
10:9	BiosRstDet[2:1]: BIOS reset detect bits[2:1]. See: BiosRstDet[0].														
8	DefSubLnk: default sublink. Read-only. Cold reset: 0. Used in conjunction with D18F0x6C[DefLnk] . 0=Sublink 0. 1=Sublink 1.														
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. BiosRstDet[2:0] = {BiosRstDet[2:1], BiosRstDet[0]}. Cold reset: 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.														
3:2	DefLnk[1:0]: default link. Read-only. DefLnk[2:0] = {DefLnk[2], DefLnk[1:0]}. Cold reset: 000b. DefLnk[2:0] is updated every time an incoming request is received with the link ID of the link on which the packet arrived. It is used by hardware to route packets during initialization, while D18F0x6C[RouteTblDisRouting]=1 , and only one outstanding request is active in the system at a time. During this interval, responses are routed to the link indicated by this field. Thus, responses are properly returned to the link, or to this node, based on the source of the request. D18F0x6C[DefSubLnk] is used to delineate sublinks as well. <table border="1" data-bbox="267 1596 941 1837"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Request came from link 0.</td> </tr> <tr> <td>001b</td> <td>Request came from link 1.</td> </tr> <tr> <td>010b</td> <td>Request came from link 2.</td> </tr> <tr> <td>011b</td> <td>Request came from link 3.</td> </tr> <tr> <td>100b</td> <td>Request came from a core on same node.</td> </tr> <tr> <td>111b-101b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	Request came from link 0.	001b	Request came from link 1.	010b	Request came from link 2.	011b	Request came from link 3.	100b	Request came from a core on same node.	111b-101b	Reserved
Bits	Description														
000b	Request came from link 0.														
001b	Request came from link 1.														
010b	Request came from link 2.														
011b	Request came from link 3.														
100b	Request came from a core on same node.														
111b-101b	Reserved														

1	ReqDis: request disable. Read-write; set-by-hardware. This bit specifies if the node is allowed to generate request packets. Reset: It resets to 0 for the BSP and to 1 for all other nodes. This bit should be cleared by BIOS once the system has been initialized from the BSP. 0=Request packets may be generated. 1=Request packets may not be generated. See 2.3 [Processor Initialization].
0	RouteTblDis: routing table disable. Read-write. Reset: 1. 1=Responses are routed based on D18F0x6C[DefLnk] and configuration-space requests received by this node are treated as if they target this node regardless of the bus number and device number. 0=Packages are routed according to D18F0x[5C:40] [Routing Table]. Once the routing tables have been set up this bit should be cleared.

D18F0x[E0,C0,A0,80] Link Capabilities

See 2.12.1.3.1 [Link Specific Registers]. This register is derived from the link capabilities register defined in the link specification. This register is 0000_0000h for link X if link X is not supported by the node.

Table 134: Register Mapping for D18F0x[E0,C0,A0,80]

Register	Function
D18F0x80	Link 0
D18F0xA0	Link 1
D18F0xC0	Link 2
D18F0xE0	Link 3

Bits	Description
31:29	CapType: capability type. Read-only. Reset: 001b.
28	DropOnUnInit: drop on uninitialized link. Read-only. Reset: 0.
27	InbndEocErr: inbound end-of-chain error. Read-only. Reset: 0.
26	ActAsSlave: act as slave. Read-only. Reset: 0.
25	Reserved.
24	HostHide. Read-only. Reset: 1.
23	ChainSide. Read-only. Reset: 0.
22:18	DevNum: device number. Read-only. Reset: 0.
17	DbIEnded: double ended. Read-only. Reset: 0.
16	WarmReset. Read-only. Reset: 1.

15:8	<p>CapPtr: capabilities pointer. Read-only. Value: Product-specific. Specifies the offset of the next link capabilities block based on the links that are supported by the node. Depending on which links are supported, this may be A0h, C0h, E0h, or 00h (in the case of the last link). The next link may not be consecutive; E.g. link 1 may point to link 3.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>This link is disabled or this is the last link; there is no next link.</td> </tr> <tr> <td>9Fh-01h</td> <td>Reserved</td> </tr> <tr> <td>A0h</td> <td>Link 1 is the next.</td> </tr> <tr> <td>BFh-A1h</td> <td>Reserved</td> </tr> <tr> <td>C0h</td> <td>Link 2 is the next.</td> </tr> <tr> <td>DFh-C1h</td> <td>Reserved</td> </tr> <tr> <td>E0h</td> <td>Link 3 is the next.</td> </tr> <tr> <td>FFh-E1h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00h	This link is disabled or this is the last link; there is no next link.	9Fh-01h	Reserved	A0h	Link 1 is the next.	BFh-A1h	Reserved	C0h	Link 2 is the next.	DFh-C1h	Reserved	E0h	Link 3 is the next.	FFh-E1h	Reserved
<u>Bits</u>	<u>Description</u>																		
00h	This link is disabled or this is the last link; there is no next link.																		
9Fh-01h	Reserved																		
A0h	Link 1 is the next.																		
BFh-A1h	Reserved																		
C0h	Link 2 is the next.																		
DFh-C1h	Reserved																		
E0h	Link 3 is the next.																		
FFh-E1h	Reserved																		
7:0	<p>CapID: capabilities ID. Read-only. Reset: 08h. Indicates HyperTransport™ technology capability.</p>																		

D18F0x[E4,C4,A4,84] Link Control

See 2.12.1.3.1 [Link Specific Registers]. This register is derived from the link control register defined in the link specification.

Table 135: Register Mapping for D18F0x[E4,C4,A4,84]

Register	Function
D18F0x84	Link 0
D18F0xA4	Link 1
D18F0xC4	Link 2
D18F0xE4	Link 3

Bits	Description										
31	Reserved.										
30:28	<p>WidthOut: link width out. Read-write; reset-applied. Specifies the operating width of the outgoing link.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>8 bits</td> </tr> <tr> <td>001b</td> <td>16 bits</td> </tr> <tr> <td>110b-010b</td> <td>Reserved</td> </tr> <tr> <td>111b</td> <td>Not connected</td> </tr> </tbody> </table> <p>The cold reset value of this field depends on the widths of the links of the connecting device, per the link specification. This field cannot be set to 16 bits when reganging a link until D18F0x[18C:170][Ganged] has been set to 1. This field cannot be changed by software if the link was determined to be disconnected by hardware at cold reset. After this field is written to by software, the link width does not change until either a warm reset or a link disconnect sequence occurs through <code>LDTSTOP_L</code>.</p>	<u>Bits</u>	<u>Description</u>	000b	8 bits	001b	16 bits	110b-010b	Reserved	111b	Not connected
<u>Bits</u>	<u>Description</u>										
000b	8 bits										
001b	16 bits										
110b-010b	Reserved										
111b	Not connected										
27	Reserved.										

26:24	WidthIn: link width in. Read-write; reset-applied. Specifies the operating width of the incoming link. See D18F0x[E4,C4,A4,84][WidthOut] for legal values. The cold reset value of this field depends on the widths of the links of the connecting device, per the link specification. This field cannot be set to 16 bits when re-gang a link until D18F0x[18C:170][Ganged] has been set to 1. This field cannot be changed by software if the link was determined to be disconnected by hardware at cold reset. After this field is written to by software, the link width does not change until either a warm reset or a link disconnect sequence occurs through <code>LDTSTOP_L</code> .
23	Reserved.
22:20	MaxWidthOut: max link width out. Read-only. Specifies the width of the outgoing link to be 8 bits or 16 bits wide, depending on the processor version. See D18F0x[E4,C4,A4,84][WidthOut] for the encoding. Indicates an 8-bit link if the link is un-ganged.
19	Reserved.
18:16	MaxWidthIn: max link width in. Read-only. Specifies the width of the incoming link to be 8 bits or 16 bits wide, depending on the processor version. See D18F0x[E4,C4,A4,84][WidthOut] for the encoding. Indicates an 8-bit link if the link is un-ganged.
15	Addr64BitEn: 64-bit address packet enable. Read-write. Cold reset: 0. 1=Requests to addresses greater than <code>FF_FFFF_FFFFh</code> are supported by this IO link. 0=Requests to addresses greater than <code>FF_FFFF_FFFFh</code> are master aborted as if the end of chain was reached. BIOS is required to ensure that the link-specification-defined “64 Bit Address Feature” bit in the device on the other side of the link is set prior to setting this bit. For coherent links, this bit is unused. D18F0x68[CHtExtAddrEn] is required to be set if this bit is set for any IO link. The link specification indicates that this bit is cleared by a warm reset; therefore this bit may be in a different state than an IO device on the other side of the link after a warm reset; care should be taken by BIOS to place devices on both sides of the link in the same state after a warm reset, before any packets to the high-order addresses enabled by this bit are generated.
14	ExtCTL: extended control time during initialization. Read-write. Cold reset: 0. This specifies the time in which the link CTL signal is held asserted during the initialization sequence that follows an <code>LDTSTOP_L</code> deassertion, after CTL is detected asserted. 0=At least 16 bit times. 1=About 50 us. ExtCTL is ignored at Gen3 frequencies.
13	LdtStopTriEn: LDTSTOP_L tristate enable. Read-write. Cold reset: 0. BIOS: <code>IF (C32r1) THEN 1 ELSE 0 ENDIF</code> . This bit is ignored by hardware when the link is operating at Gen3 frequencies. 1=During the <code>LDTSTOP_L</code> disconnect sequence, the link transmitter signals are placed into the high-impedance state and the receivers are prepared for the high-impedance mode. For the receivers, this includes cutting power to the receiver differential amplifiers and ensuring that there are no resultant high-current paths in the circuits. 0=During the <code>LDTSTOP_L</code> disconnect sequence, the link transmitter signals are driven, but in an undefined state, and the link receiver signals are assumed to be undriven.
12	IsocEn: isochronous flow-control mode enable. Read-write; reset-applied. Cold reset: 0. 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. All coherent links of the system must use the same flow-control mode. See 2.9.3.2.5 [Display Refresh And IFCM] .
11:10	Reserved.

9:8	CrcErr: CRC Error. Read; set-by-hardware; write-1-to-clear. Cold reset: 00b. Bit[1] applies to the upper byte of the link and bit[0] applies to the lower byte. 1=The hardware detected a CRC error on the incoming link while not in retry mode; if in retry mode, then bit[8] may be set to indicate an uncorrectable error was detected; such uncorrectable error cases are: Link reconnect fails exceeding the limit in D18F0x150 [Link Global Retry Control][TotalRetryAttempts] .
7	TransOff: transmitter off. Read-write. 1=Output tristates. Notes about TransOff and EndOfChain: The initial state is programmed by hardware after each cold reset (0 if the link connects; 1 if it does not connect; see 2.12.1.3 [Link Type Detect]). If D18F0x16C[ConnDly]=0 : setting these bits causes the link behavior to change immediately; clearing these bits is illegal; warm reset and LDTSTOP_L do not affect the state of these bits; TransOff may only be set if EndOfChain is set. If D18F0x16C[ConnDly]=1 : values written to these bits take effect during the next warm reset or LDTSTOP_L. Reads represent the last value written rather than the current link behavior; TransOff and EndOfChain should always be programmed to the same state. BIOS: For dual-node processors, BIOS should program this according to 2.12.1.5 [Link Mapping Between Package and Node] .
6	EndOfChain: end of chain. Read-write. 1=The link is not part of the logical HyperTransport™ technology chain; packets targeting this link are dropped. See TransOff, above. BIOS: For dual-node processors, BIOS should program this according to 2.12.1.5 [Link Mapping Between Package and Node] .
5	InitComplete: initialization complete. Read-only; set-by-hardware. Reset: 0. This bit is set by hardware when low-level link initialization has successfully completed. If there is no device on the other end of the link, or if the device on the other side of the link is unable to properly perform link initialization, then the bit is not set. This bit is not cleared for LDTSTOP_L disconnects or retries. Hardware may report 0 during BIST mode or ILM.
4	LinkFail: link failure. Read; set-by-hardware; write-1-to-clear. Cold reset: 0. This bit is set high by the hardware when a CRC error is detected on the link (if enabled by CrcFloodEn), the link fails to reconnect, if a sync flood is received by the link, or if the link is not used in the system. See 2.13.1.9.1 [Common Diagnosis Information] .
3	CrcForceErr: CRC force error command. Read-write. Reset: 0. 1=The link transmission logic generates erroneous periodic or per-packet CRC values on all enabled byte lanes. 0=Transmitted CRC values match the values calculated per the link specification. This bit is intended to be used to check the CRC failure detection logic of the device on the other side of the link. See D18F0x150[ForceErrType] for retry mode.
2	Reserved.
1	CrcFloodEn: CRC flood enable. Read-write. Reset: 0. 1=Setting either of the CrcErr bits results in sync packets to all enabled outgoing links and the D18F0x[E4,C4,A4,84][LinkFail] bit is set. 0=Setting either of the CrcErr bits do not result in sync packets or setting the D18F0x[E4,C4,A4,84][LinkFail] bit. In Gen3 protocol, exceeding the D18F0x150[TotalRetryAttempts] limit results in a sync flood regardless of how CrcFloodEn is set. The resulting sync flood does not propagate to other links or set Linkfail unless CrcFloodEn is set. This bit is ignored if D18F3x44[SyncPktGenDis] is set.
0	Reserved.

[D18F0x\[E8,C8,A8,88\] Link Frequency/Revision](#)

See [2.12.1.3.1 \[Link Specific Registers\]](#). This register is derived from the link frequency/revision register

defined in the link specification.

Table 136: Register Mapping for D18F0x[E8,C8,A8,88]

Register	Function
D18F0x88	Link 0
D18F0xA8	Link 1
D18F0xC8	Link 2
D18F0xE8	Link 3

Bits	Description																																																																				
31:16	<p>LnkFreqCap: link frequency capability. Read-only. LnkFreqCap[30:0] = {D18F0x[FC,DC,BC,9C][FreqCapExt[14:0]], LnkFreqCap[15:0]}. Reset: Product-specific. 1=The link frequency is supported; 0=The link frequency is not supported. Indicates logical support for these frequencies; however, electrical support for these frequencies may vary based on the part number and other system considerations.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>0.2 GHz. (this bit is 1 in all products).</td> <td>[16]</td> <td>2.8 GHz.</td> </tr> <tr> <td>[1]</td> <td>0.3 GHz. (this bit is 0 in all products).</td> <td>[17]</td> <td>3.0 GHz.</td> </tr> <tr> <td>[2]</td> <td>0.4 GHz.</td> <td>[18]</td> <td>3.2 GHz.</td> </tr> <tr> <td>[3]</td> <td>0.5 GHz. (this bit is 0 in all products).</td> <td>[19]</td> <td>Reserved</td> </tr> <tr> <td>[4]</td> <td>0.6 GHz.</td> <td>[20]</td> <td>Reserved</td> </tr> <tr> <td>[5]</td> <td>0.8 GHz.</td> <td>[30:21]</td> <td>Reserved</td> </tr> <tr> <td>[6]</td> <td>1.0 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[7]</td> <td>1.2 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[8]</td> <td>1.4 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[9]</td> <td>1.6 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[10]</td> <td>1.8 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[11]</td> <td>2.0 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[12]</td> <td>2.2 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[13]</td> <td>2.4 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[14]</td> <td>2.6 GHz.</td> <td></td> <td></td> </tr> <tr> <td>[15]</td> <td>Reserved.</td> <td></td> <td></td> </tr> </tbody> </table>	Bit	Description	Bit	Description	[0]	0.2 GHz. (this bit is 1 in all products).	[16]	2.8 GHz.	[1]	0.3 GHz. (this bit is 0 in all products).	[17]	3.0 GHz.	[2]	0.4 GHz.	[18]	3.2 GHz.	[3]	0.5 GHz. (this bit is 0 in all products).	[19]	Reserved	[4]	0.6 GHz.	[20]	Reserved	[5]	0.8 GHz.	[30:21]	Reserved	[6]	1.0 GHz.			[7]	1.2 GHz.			[8]	1.4 GHz.			[9]	1.6 GHz.			[10]	1.8 GHz.			[11]	2.0 GHz.			[12]	2.2 GHz.			[13]	2.4 GHz.			[14]	2.6 GHz.			[15]	Reserved.		
Bit	Description	Bit	Description																																																																		
[0]	0.2 GHz. (this bit is 1 in all products).	[16]	2.8 GHz.																																																																		
[1]	0.3 GHz. (this bit is 0 in all products).	[17]	3.0 GHz.																																																																		
[2]	0.4 GHz.	[18]	3.2 GHz.																																																																		
[3]	0.5 GHz. (this bit is 0 in all products).	[19]	Reserved																																																																		
[4]	0.6 GHz.	[20]	Reserved																																																																		
[5]	0.8 GHz.	[30:21]	Reserved																																																																		
[6]	1.0 GHz.																																																																				
[7]	1.2 GHz.																																																																				
[8]	1.4 GHz.																																																																				
[9]	1.6 GHz.																																																																				
[10]	1.8 GHz.																																																																				
[11]	2.0 GHz.																																																																				
[12]	2.2 GHz.																																																																				
[13]	2.4 GHz.																																																																				
[14]	2.6 GHz.																																																																				
[15]	Reserved.																																																																				
15:12	Reserved.																																																																				

11:8	<p>Freq[3:0]: link frequency. Read-write; Reset-applied. Freq[4:0] = {D18F0x[FC,DC,BC,9C][Freq[4]], D18F0x[E8,C8,A8,88][Freq[3:0]]}. Cold reset: 0h. When D18F0x[FC,DC,BC,9C][Freq[4]] is set or cleared, D18F0x[E8,C8,A8,88][Freq[3:0]] must be written to the correct value before D18F0x[FC,DC,BC,9C][Freq[4]] is written. After this field is updated, the link frequency does not change until either a warm reset or a link disconnect sequence occurs through LDTSTOP_L. The value read from this field is the last value written. Writes to this field are ignored if a non-supported frequency is written.</p> <table border="1" data-bbox="264 436 1144 1045"> <thead> <tr> <th>Bits</th> <th>Link Frequency</th> <th>Bits</th> <th>Link Frequency</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0.2 GHz.²</td> <td>10h</td> <td>Reserved</td> </tr> <tr> <td>01h</td> <td>Reserved</td> <td>11h</td> <td>2.8 GHz.</td> </tr> <tr> <td>02h</td> <td>0.4 GHz.²</td> <td>12h</td> <td>3.0 GHz.</td> </tr> <tr> <td>03h</td> <td>Reserved</td> <td>13h</td> <td>3.2 GHz.</td> </tr> <tr> <td>04h</td> <td>0.6 GHz.²</td> <td>14h</td> <td>Reserved</td> </tr> <tr> <td>05h</td> <td>0.8 GHz.²</td> <td>15h</td> <td>Reserved</td> </tr> <tr> <td>06h</td> <td>1.0 GHz.^{1,2}</td> <td>1Fh-16h</td> <td>Reserved</td> </tr> <tr> <td>07h</td> <td>1.2 GHz.</td> <td></td> <td></td> </tr> <tr> <td>08h</td> <td>1.4 GHz.</td> <td></td> <td></td> </tr> <tr> <td>09h</td> <td>1.6 GHz.</td> <td></td> <td></td> </tr> <tr> <td>0Ah</td> <td>1.8 GHz.</td> <td></td> <td></td> </tr> <tr> <td>0Bh</td> <td>2.0 GHz.</td> <td></td> <td></td> </tr> <tr> <td>0Ch</td> <td>2.2 GHz.²</td> <td></td> <td></td> </tr> <tr> <td>0Dh</td> <td>2.4 GHz.</td> <td></td> <td></td> </tr> <tr> <td>0Eh</td> <td>2.6 GHz.²</td> <td></td> <td></td> </tr> <tr> <td>0Fh</td> <td>Reserved.</td> <td></td> <td></td> </tr> </tbody> </table> <p>Note:</p> <ol style="list-style-type: none"> G34r1: 1.0 GHz disallowed for links connecting internal nodes as indicated by D18F0x1A0[IntLnkRoute]. AM3r2: 0.4-1.0 GHz disallowed for AM3r2 package. 0.2 GHz supported only as boot frequency to first warm reset; 0.2 GHz disallowed as operational frequency. 	Bits	Link Frequency	Bits	Link Frequency	00h	0.2 GHz. ²	10h	Reserved	01h	Reserved	11h	2.8 GHz.	02h	0.4 GHz. ²	12h	3.0 GHz.	03h	Reserved	13h	3.2 GHz.	04h	0.6 GHz. ²	14h	Reserved	05h	0.8 GHz. ²	15h	Reserved	06h	1.0 GHz. ^{1,2}	1Fh-16h	Reserved	07h	1.2 GHz.			08h	1.4 GHz.			09h	1.6 GHz.			0Ah	1.8 GHz.			0Bh	2.0 GHz.			0Ch	2.2 GHz. ²			0Dh	2.4 GHz.			0Eh	2.6 GHz. ²			0Fh	Reserved.		
Bits	Link Frequency	Bits	Link Frequency																																																																		
00h	0.2 GHz. ²	10h	Reserved																																																																		
01h	Reserved	11h	2.8 GHz.																																																																		
02h	0.4 GHz. ²	12h	3.0 GHz.																																																																		
03h	Reserved	13h	3.2 GHz.																																																																		
04h	0.6 GHz. ²	14h	Reserved																																																																		
05h	0.8 GHz. ²	15h	Reserved																																																																		
06h	1.0 GHz. ^{1,2}	1Fh-16h	Reserved																																																																		
07h	1.2 GHz.																																																																				
08h	1.4 GHz.																																																																				
09h	1.6 GHz.																																																																				
0Ah	1.8 GHz.																																																																				
0Bh	2.0 GHz.																																																																				
0Ch	2.2 GHz. ²																																																																				
0Dh	2.4 GHz.																																																																				
0Eh	2.6 GHz. ²																																																																				
0Fh	Reserved.																																																																				
7:0	<p>Revision. Read-only. Reset: 60h. Indicates that the processor is designed to version 3.00 of the link specification.</p>																																																																				

D18F0x[EC,CC,AC,8C] Link Feature Capability

See 2.12.1.3.1 [Link Specific Registers]. This register is derived from the link feature capability register defined in the link specification. Unless otherwise specified: 0=The feature is not supported; 1=The feature is supported.

Table 137: Register Mapping for D18F0x[EC,CC,AC,8C]

Register	Function
D18F0x8C	Link 0
D18F0xAC	Link 1
D18F0xCC	Link 2
D18F0xEC	Link 3

Bits	Description
31:10	Reserved.
9	UpstrCfgCap: upstream configuration capable. Read-only. Reset: 0.
8	ExtRegSet: extended register set. Read-only. Reset: 0.
7:6	Reserved.
5	UnitIdReOrderDis: UnitID reorder disable. Read-write. Read-only. Reset: 0. 1=Upstream reordering for different UnitIDs is not supported; i.e., all upstream packets are ordered as if they have the same UnitID. 0=Reordering based on UnitID is supported.
4	64BitAddr: 64-bit link addressing. Read-only. Reset: 1.
3	ExtCTLRqd: extended CTL required. Read-only. Reset: 0.
2	CrcTstMode: CRC test mode. Read-only. Reset: 0.
1	LdtStopMode: LDTSTOP_L supported. Read-only. Reset: 1.
0	IsocMode: isochronous flow control mode. Read-only. Reset: 1.

D18F0x[F0,D0,B0,90] Link Base Channel Buffer Count

Read-write; Reset-applied.

Table 138: Register Mapping for D18F0x[F0,D0,B0,90]

Register	Function
D18F0x90	Link 0
D18F0xB0	Link 1
D18F0xD0	Link 2
D18F0xF0	Link 3

D18F0x[F0,D0,B0,90] and D18F0x[F4,D4,B4,94] specify the *hard-allocated* link flow-control buffer counts in each virtual channel available to the transmitter at the other end of the link; it also provides the *free buffers* that may be used by any of the virtual channels, as needed. Base channel buffers are specified in D18F0x[F0,D0,B0,90]; isochronous buffer counts (if in IFCM) are specified in D18F0x[F4,D4,B4,94]. For all fields that specify buffer counts in D18F0x[F0,D0,B0,90] and D18F0x[F4,D4,B4,94], if the link is ganged, then the number of buffers allocated is 2 times the value of the field; If the link is unganged, then the number of buffers allocated is the value of the field. See 2.12.1.3.1 [Link Specific Registers]. See section 2.12.1.3.2 [Unused Links].

The cold or warm reset value is determined by whether the link initializes, whether the link is IO/coherent, whether the link is ganged/unganged, and whether the settings are locked by LockBc.

```
IF (LinkConnected & IoLink & LockBc) THEN Cold reset: 0285_0292h.
ELSEIF (LinkConnected & IoLink & ~LockBc) THEN Reset: 0285_0292h.
ELSEIF (LinkConnected & ~IoLink & LockBc) THEN Cold reset: 028A_9944h.
ELSEIF (LinkConnected & ~IoLink & ~LockBc) THEN Reset: 028A_9944h.
ELSE Reset: X. ENDIF. //Link not connected
```

The hard-allocated buffer counts are transmitted to the device at the other end of the link in buffer release messages after link initialization. The remaining buffers are held in the free list (specified by FreeData and FreeCmd) 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 be immediately sent back to the device at the other end of the link 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 (64 buffers) if ganged or 32 (32 buffers) if unganged:
 - $(D18F0x[F0,D0,B0,90][NpReqCmd] + D18F0x[F0,D0,B0,90][PReq] + D18F0x[F0,D0,B0,90][RspCmd] + D18F0x[F0,D0,B0,90][ProbeCmd] + D18F0x[F0,D0,B0,90][FreeCmd] + D18F0x[F4,D4,B4,94][IsocNpReqCmd] + D18F0x[F4,D4,B4,94][IsocPReq] + D18F0x[F4,D4,B4,94][IsocRspCmd]) \leq 32$.
- The total number of data buffers allocated in the base and isochronous registers of a link cannot exceed 8 (16 buffers) if ganged or 8 (8 buffers) if unganged:
 - $(D18F0x[F0,D0,B0,90][NpReqData] + D18F0x[F0,D0,B0,90][RspData] + D18F0x[F0,D0,B0,90][PReq] + D18F0x[F0,D0,B0,90][FreeData] + D18F0x[F4,D4,B4,94][IsocPReq] + D18F0x[F4,D4,B4,94][IsocNpReqData] + D18F0x[F4,D4,B4,94][IsocRspData]) \leq 8$.
- The total number of hard allocated command buffers cannot exceed 24 (48 buffers) if ganged or 24 (24 buffers) if unganged.
 - $(D18F0x[F0,D0,B0,90][ProbeCmd] + D18F0x[F0,D0,B0,90][RspCmd] + D18F0x[F0,D0,B0,90][PReq] + D18F0x[F0,D0,B0,90][NpReqCmd] + D18F0x[F4,D4,B4,94][IsocRspCmd] + D18F0x[F4,D4,B4,94][IsocRspCmd] + D18F0x[F4,D4,B4,94][IsocPReq] + D18F0x[F4,D4,B4,94][IsocNpReqCmd]) \leq 24$.
- BIOS must set up non-zero counts (and adjust the base channel counts accordingly) prior to enabling IFCM.
- If an IOMMU is present in the system, $D18F0x[F4,D4,B4,94][IsocNpReqCmd]$ must be non-zero for all enabled links.

See 2.9.3.2.5.1 [Recommended Buffer Count Settings Overview].

Table 139: BIOS Recommendations for $D18F0x[F0,D0,B0,90]$, $D18F0x[F4,D4,B4,94]$

Condition			$D18F0x[F4,D4,B4,94]$					$D18F0x[F0,D0,B0,90]$							
			IsocRspData	IsocNpReqData	IsocRspCmd	IsocPReq	IsocNpReqCmd	FreeData	FreeCmd	RspData	NpReqData	ProbeCmd	RspCmd	PReq	NpReqCmd
~IoLink	LinkGang	PrbFltrEn	0	0	0	0	1	0	8	3	3	4	9	2	8
		~PrbFltrEn	0	0	0	0	1	0	8	3	3	8	9	2	4
~IoLink	~LinkGang	PrbFltrEn	0	0	0	0	1	0	8	3	3	4	9	2	8
		~PrbFltrEn	0	0	0	0	1	0	8	3	3	8	9	2	4
IoLink	LinkGang		0	0	0	0	1	0	8	1	0 ¹	0	2	7 ¹	14 ¹

1. This recommendation is for AMD chipsets which don't generate requests in the NpReqData channel. If a non-AMD chipset generates requests in the NpReqData channel then NpReqData=1, PReq=6, NpReqCmd=15.

Bits	Description
31	LockBc: lock buffer count register. Cold reset: 0. 1=The buffer count registers, D18F0x[F0,D0,B0,90] and D18F0x[F4,D4,B4,94] 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. BIOS: Table 139 .
24:20	FreeCmd: free command buffer count. BIOS: Table 139 .
19:18	RspData: response data buffer count. BIOS: Table 139 .
17:16	NpReqData: non-posted request data buffer count. BIOS: Table 139 .
15:12	ProbeCmd: probe command buffer count. BIOS: Table 139 .
11:8	RspCmd: response command buffer count. BIOS: Table 139 .
7:5	PReq: posted request command and data buffer count. Specifies the number of posted command and posted data buffers allocated. BIOS: Table 139 .
4:0	NpReqCmd: non-posted request command buffer count. BIOS: Table 139 .

D18F0x[F4,D4,B4,94] Link Isochronous Channel Buffer Count

Read-write; Reset-applied. See [D18F0x\[F0,D0,B0,90\]](#).

Table 140: [Register Mapping](#) for [D18F0x\[F4,D4,B4,94\]](#)

Register	Function
D18F0x94	Link 0
D18F0xB4	Link 1
D18F0xD4	Link 2
D18F0xF4	Link 3

The cold or warm reset is determined by whether the link initializes and whether the settings are locked by LockBc.

```
IF (LinkConnected && D18F0x[F0,D0,B0,90][LockBc]) THEN Cold reset: 0000_0000h.
ELSEIF (LinkConnected && ~D18F0x[F0,D0,B0,90][LockBc]) THEN Reset: 0000_0000h.
ELSE Reset: X. ENDIF. //Link not connected
```

Bits	Description
31:29	Reserved.
28:27	IsocRspData: isochronous response data buffer count. BIOS: Table 139 .
26:25	IsocNpReqData: isochronous non-posted request data buffer count. BIOS: Table 139 .
24:22	IsocRspCmd: isochronous response command buffer count. BIOS: Table 139 .
21:19	IsocPReq: isochronous posted request command and data buffer count. BIOS: Table 139 . This specifies the number of isochronous posted command and posted data buffers allocated.
18:16	IsocNpReqCmd: isochronous non-posted request command buffer count. BIOS: Table 139 .

15:8	SecBusNum: secondary bus number. 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.

D18F0x[F8,D8,B8,98] Link Type

See [2.12.1.3.1 \[Link Specific Registers\]](#).

Table 141: [Register Mapping](#) for [D18F0x\[F8,D8,B8,98\]](#)

Register	Function
D18F0x98	Link 0
D18F0xB8	Link 1
D18F0xD8	Link 2
D18F0xF8	Link 3

Bits	Description
31:5	Reserved.
4	LinkConPend: link connect pending. Read-only. Reset: X. 1=Hardware is currently determining if the link is connected to another device. 0=The link connection has been determined. This bit qualifies the LinkCon bit.
3	Reserved.
2	NC: non coherent. Read-only. Reset: X. This bit specifies the link type. 0=coherent link. 1=IO link.
1	InitComplete: initialization complete. Read-only. Reset: X. 1=Link initialization is complete. This is a duplicate of D18F0x[E4,C4,A4,84] [Link Control][InitComplete] . The NC bit is invalid until link initialization is complete.
0	LinkCon: link connected. Read-only. Reset: X. 1=The link is connected to another device. 0=The link is not connected. This is not valid until LinkConPend=0.

D18F0x[FC,DC,BC,9C] Link Frequency Extension

See [2.12.1.3.1 \[Link Specific Registers\]](#). This register is derived from link specification.

Table 142: [Register Mapping](#) for [D18F0x\[FC,DC,BC,9C\]](#)

Register	Function
D18F0x9C	Link 0
D18F0xBC	Link 1
D18F0xDC	Link 2
D18F0xFC	Link 3

Bits	Description
31:16	Reserved.

15:1	FreqCapExt: link frequency capability extension. See: D18F0x[E8,C8,A8,88][LnkFreqCap] .
0	Freq[4]: link frequency. See: D18F0x[E8,C8,A8,88][Freq[3:0]] .

D18F0x[11C,118,114,110] Link Clumping Enable

Reset: 0000_0000h. [D18F0x\[11C,118,114,110\]](#) are associated with the whole link if it is ganged or sublink 0 if it is unganged; [D18F0x\[12C,128,124,120\]](#) are associated with sublink 1 if the link is unganged. If the node does not support a link, then the corresponding register addresses become reserved.

Table 143: [Register Mapping](#) for [D18F0x\[11C,118,114,110\]](#)

Register	Function
D18F0x110	Link 0
D18F0x114	Link 1
D18F0x118	Link 2
D18F0x11C	Link 3

These registers specify 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.

D18F0x[12C,128,124,120] Sublink 1 Clumping Enable

Reset: 0000_0000h. See [D18F0x\[11C,118,114,110\]](#).

Table 144: [Register Mapping](#) for [D18F0x\[12C,128,124,120\]](#)

Register	Function
D18F0x120	Link 0
D18F0x124	Link 1
D18F0x128	Link 2
D18F0x12C	Link 3

Bits	Description
31:0	See: D18F0x[11C,118,114,110] . If a link is ganged or not supported, then the corresponding register in this group is reserved.

D18F0x[14C:130] Link Retry

Table 145: Register Mapping for D18F0x[14C:130]

Register	Function
D18F0x130	Link 0 Sublink 0
D18F0x134	Link 1 Sublink 0
D18F0x138	Link 2 Sublink 0
D18F0x13C	Link 3 Sublink 0
D18F0x140	Link 0 Sublink 1
D18F0x144	Link 1 Sublink 1
D18F0x148	Link 2 Sublink 1
D18F0x14C	Link 3 Sublink 1

If a link is ganged, then the sublink 0 retry register specifies the link retry function and the sublink 1 retry register is reserved. If a link is not supported by the node, then both the sublink 0 and sublink 1 retry registers are reserved. These registers are reserved if [D18F3xE8\[LnkRtryCap\]](#)=0.

Bits	Description
31:16	RetryCount. Read-write; updated-by-hardware. Cold reset: 0. This is a 16-bit counter that is incremented by hardware. The counter is incremented in two ways, (1) the counter increments once for each failed training attempt and (2) the counter increments once for each packet error that causes a retry attempt. If the counter value is FFFFh it increments to 0000h and the RetryCountRollover bit is set. RetryCount is not incremented for retries initiated by other devices, only for errors detected by the node.
15:13	Reserved.
12	DataCorruptOut: sent corrupted data. Read; write-1-to-clear; set-by-hardware. Cold reset: 0. 1=Data sent on the link was marked with Data Error to indicate that it is known to be corrupted.
11	InitFail. Read; write-1-to-clear; set-by-hardware. Cold reset: 0. 1=Initialization sequence failed on a link reconnect.
10	StompedPktDet: stomped packet detected by receiver. Read; write-1-to-clear; set-by-hardware. Cold reset: 0.
9	RetryCountRollover. Read; write-1-to-clear; set-by-hardware. Cold reset: 0. See RetryCount .
8	RetryErrorDet: retry error detected. Read; write-1-to-clear; set-by-hardware. Cold reset: 0. 1=A retry was initiated in one of the ways listed in RetryCount .
7:6	ShortRetryAttempts. Read-write. Reset: 11b. This specifies the number of short retry attempts when operating at a Gen3 link frequency; after exceeding this value, long retries are attempted until the max count specified by D18F0x150 [Link Global Retry Control][TotalRetryAttempts] is exceeded. The retry attempt counter is not incremented for retries initiated by other devices, only for errors detected by the node. This field is ignored when operating at Gen1 link frequencies.
5:4	Reserved.
3	DisRetryDataError: disable link retry on data packet error. Read-write. Reset: 0. 1=The node does not initiate the retry sequence if an error is detected on a data packet; Data packets are acknowledged even if there is a CRC error. This is intended to support debug modes in which errors are detected but allowed to propagate through the crossbar in order to allow logging of error data patterns in trace mode.

2	DisRetryAnyError: disable link retry on any packet error. Read-write. Reset: 0. 1=The node does not initiate the retry sequence if an error is detected; Packets are acknowledged even if there is a CRC error. This is intended to support debug modes in which errors are detected but allowed to propagate through the crossbar in order to allow logging of error data patterns in trace mode.
1	ForceRetryError. Read-write; cleared-by-hardware. Cleared by hardware once the error has been injected onto the link. Reset: 0. This bit may be used by diagnostic software to test the error detection and retry logic of the link. 1=Forces a CRC error in one packet from the transmitter. See D18F0x150 [Link Global Retry Control] [MultRetryErr].
0	RetryModeEnable. Read-write; reset-applied. Cold reset: 0. 1=Place the link in error retry mode when reconnecting after the next warm reset. See Table 117 .

D18F0x150 Link Global Retry Control

This register is reserved if [D18F3xE8\[LnkRtryCap\]](#)=0. All fields of this register are expected to be programmed the same in all nodes of the system (except ForceErrType and MultRetryErr).

Bits	Description
31:19	Reserved.
18:16	TotalRetryAttempts. Read-write. Cold reset: 111b. BIOS: This register should be programmed to values of 1 or greater. Specifies the total number of retry attempts (short and long) allowed on any link before the link is considered to have failed. When operating at Gen3 link frequencies, short retry attempts are limited by D18F0x[14C:130] [Link Retry] [ShortRetryAttempts]; the remaining are long retry attempts. The link is determined to have failed after TotalRetryAttempts + 1 errors; e.g., if TotalRetryAttempts=7, then the link is determined to have failed as a result of the 8 errors. The retry attempt counter for a link is incremented each time D18F0x[14C:130] [RetryCount] for that link is incremented.
15:14	Reserved.
13	HtRetryCrcDatInsDynEn: link retry CRC data insertion enable. Read-write. Cold reset: 0. BIOS: 1. 1=Enables dynamic mode for CRC insertion in data packets on a coherent link. In this mode, the transmitter follows the insertion policy defined by HtRetryCrcDatIns[2:0] for a link which is close to idle; however, it inserts fewer CRC cells as the link becomes busy. Must be modified only when all links are operating at Gen1 frequency.
12	HtRetryCrcCmdPackDynEn: link retry CRC command packet dynamic mode enable. Read-write. Cold reset: 0. BIOS: 1. 1=Enables dynamic mode for CRC command packing on a coherent link with retry enabled. In this mode, command packing is suspended when a link is not busy. This field is valid only when HtRetryCrcCmdPack =1. Must be modified only when all links are operating at Gen1 frequency.

11:9	<p>HtRetryCrcDatIns: link retry CRC data insertion. Read-write. Cold reset: 000b. BIOS: 100b. Specifies insertion of additional CRC cells in a data packet over coherent link. A data packet is defined as a data command header followed by at most 4 data beats (beat 0 through beat 3) of 16 bytes each with a data packet CRC at the end. This field must be modified only when all links are operating at Gen1 frequency.</p> <p>IF (<u>LinkFreq</u>≤<u>NBCOF</u>) THEN</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>No additional CRC insertion</td> </tr> <tr> <td>001b</td> <td>(<u>PROC</u>≥<u>OR_C0</u>): CRC insertion after data beat 1 (<u>PROC</u><<u>OR_C0</u>): CRC insertion after data beat 0</td> </tr> <tr> <td>010b</td> <td>(<u>PROC</u>≥<u>OR_C0</u>): CRC insertion after cmd header and after data beat 1 (<u>PROC</u><<u>OR_C0</u>): CRC insertion after cmd header and after data beat 0</td> </tr> <tr> <td>011b</td> <td>CRC insertion after cmd header, data beat 0 and data beat 1</td> </tr> <tr> <td>100b</td> <td>CRC insertion after cmd header, data beat 0, data beat 1 and data beat 2</td> </tr> <tr> <td>111b-101b</td> <td>Reserved</td> </tr> </tbody> </table> <p>ELSE // (<u>LinkFreq</u>><u>NBCOF</u>)</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>No additional CRC insertion</td> </tr> <tr> <td>100b-001b</td> <td>(<u>PROC</u>≥<u>OR_C0</u>): CRC insertion after cmd header and after data beat 1 (<u>PROC</u><<u>OR_C0</u>): CRC insertion after cmd header and after data beat 0</td> </tr> <tr> <td>111b-101b</td> <td>Reserved</td> </tr> </tbody> </table> <p>ELSE.</p>	<u>Bits</u>	<u>Description</u>	000b	No additional CRC insertion	001b	(<u>PROC</u> ≥ <u>OR_C0</u>): CRC insertion after data beat 1 (<u>PROC</u> < <u>OR_C0</u>): CRC insertion after data beat 0	010b	(<u>PROC</u> ≥ <u>OR_C0</u>): CRC insertion after cmd header and after data beat 1 (<u>PROC</u> < <u>OR_C0</u>): CRC insertion after cmd header and after data beat 0	011b	CRC insertion after cmd header, data beat 0 and data beat 1	100b	CRC insertion after cmd header, data beat 0, data beat 1 and data beat 2	111b-101b	Reserved	<u>Bits</u>	<u>Description</u>	000b	No additional CRC insertion	100b-001b	(<u>PROC</u> ≥ <u>OR_C0</u>): CRC insertion after cmd header and after data beat 1 (<u>PROC</u> < <u>OR_C0</u>): CRC insertion after cmd header and after data beat 0	111b-101b	Reserved
<u>Bits</u>	<u>Description</u>																						
000b	No additional CRC insertion																						
001b	(<u>PROC</u> ≥ <u>OR_C0</u>): CRC insertion after data beat 1 (<u>PROC</u> < <u>OR_C0</u>): CRC insertion after data beat 0																						
010b	(<u>PROC</u> ≥ <u>OR_C0</u>): CRC insertion after cmd header and after data beat 1 (<u>PROC</u> < <u>OR_C0</u>): CRC insertion after cmd header and after data beat 0																						
011b	CRC insertion after cmd header, data beat 0 and data beat 1																						
100b	CRC insertion after cmd header, data beat 0, data beat 1 and data beat 2																						
111b-101b	Reserved																						
<u>Bits</u>	<u>Description</u>																						
000b	No additional CRC insertion																						
100b-001b	(<u>PROC</u> ≥ <u>OR_C0</u>): CRC insertion after cmd header and after data beat 1 (<u>PROC</u> < <u>OR_C0</u>): CRC insertion after cmd header and after data beat 0																						
111b-101b	Reserved																						
8	<p>HtRetryCrcCmdPack: link retry CRC command packing. Read-write. Cold reset: 0. BIOS: 1. 1=Enables command packing on coherent links with retry enabled. Command packing allows a coherent link transmitter to pack multiple commands together with a single CRC. Must be modified only when all links are operating at Gen1 frequency.</p>																						
7	Reserved.																						
6:5	<p>ForceErrType: force error type. Read-write. Cold reset: 00b. Specifies the error type generated by D18F0x[14C:130][ForceRetryError], D18F0x[E4,C4,A4,84][CrcForceErr], and D18F3x44[GenCrcErrByte1, GenCrcErrByte0].</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Forces per-packet CRC error in any packet type (NOP, command, or data).</td> </tr> <tr> <td>01b</td> <td>Forces per-packet CRC error on a command packet only (not including NOP).</td> </tr> <tr> <td>10b</td> <td>Forces per-packet CRC error on a data packet only. If <u>HtRetryCrcDatIns</u>=1, then the error is forced into the first CRC of the packet.</td> </tr> <tr> <td>11b</td> <td>Forces per-packet CRC error on a data packet only. If <u>HtRetryCrcDatIns</u>=1, then the error is forced into the last CRC of the packet.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Forces per-packet CRC error in any packet type (NOP, command, or data).	01b	Forces per-packet CRC error on a command packet only (not including NOP).	10b	Forces per-packet CRC error on a data packet only. If <u>HtRetryCrcDatIns</u> =1, then the error is forced into the first CRC of the packet.	11b	Forces per-packet CRC error on a data packet only. If <u>HtRetryCrcDatIns</u> =1, then the error is forced into the last CRC of the packet.												
<u>Bits</u>	<u>Description</u>																						
00b	Forces per-packet CRC error in any packet type (NOP, command, or data).																						
01b	Forces per-packet CRC error on a command packet only (not including NOP).																						
10b	Forces per-packet CRC error on a data packet only. If <u>HtRetryCrcDatIns</u> =1, then the error is forced into the first CRC of the packet.																						
11b	Forces per-packet CRC error on a data packet only. If <u>HtRetryCrcDatIns</u> =1, then the error is forced into the last CRC of the packet.																						
4	<p>MultRetryErr: multiple retry force error. Read-write. Cold reset: 0. 1=Inhibits hardware clearing of D18F0x[14C:130][Link Retry][ForceRetryError], thereby causing multiple link retry errors (at a very high rate). This can be used to test software associated with reporting of multiple link reconnect failures.</p>																						
3:0	Reserved.																						

D18F0x160 Extended Node ID

Bits	Description
31:19	Reserved.

18:16	CpuCnt[7:5]: CPU count bits[7:5]. See: D18F0x60[CpuCnt[4:0]] .
15:0	Reserved.

D18F0x164 Coherent Link Traffic Distribution

Reset: 0000_0000h. See [2.9.3.2.4 \[Link Traffic Distribution\]](#) for details about link traffic distribution.

Bits	Description																				
31:24	Reserved.																				
23:16	<p>DstLnk: distribution destination link. Read-write. Specifies the pool of links over which traffic is distributed. Packets which are not eligible for distribution (for example sized reads and writes) are routed normally, based on the routing tables. If the link is ganged, then only the sublink 0 bit needs to be set; the sublink 1 bit is ignored. See cHTVicDistMode.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>link 0, sublink 0</td> <td>[4]</td> <td>link 0, sublink 1</td> </tr> <tr> <td>[1]</td> <td>link 1, sublink 0</td> <td>[5]</td> <td>link 1, sublink 1</td> </tr> <tr> <td>[2]</td> <td>link 2, sublink 0</td> <td>[6]</td> <td>link 2, sublink 1</td> </tr> <tr> <td>[3]</td> <td>link 3, sublink 0</td> <td>[7]</td> <td>link 3, sublink 1</td> </tr> </tbody> </table>	Bit	Description	Bit	Description	[0]	link 0, sublink 0	[4]	link 0, sublink 1	[1]	link 1, sublink 0	[5]	link 1, sublink 1	[2]	link 2, sublink 0	[6]	link 2, sublink 1	[3]	link 3, sublink 0	[7]	link 3, sublink 1
Bit	Description	Bit	Description																		
[0]	link 0, sublink 0	[4]	link 0, sublink 1																		
[1]	link 1, sublink 0	[5]	link 1, sublink 1																		
[2]	link 2, sublink 0	[6]	link 2, sublink 1																		
[3]	link 3, sublink 0	[7]	link 3, sublink 1																		
15:11	Reserved.																				
10:8	<p>DstNode: coherent link distribution destination node. Read-write. For cHTReqDistEn and cHTRspDistEn, DstNode[2:0] specifies the destination node for which coherent link traffic should be distributed. For cHTPrbDistEn, all probes originating from the local node are distributed irrespective of the value of DstNode[2:0]. Packets specified by cHTReqDistEn, cHTRspDistEn, and cHTPrbDistEn that are destined for DstNode are distributed between links specified by DstLnk in approximately a round-robin fashion.</p>																				
7:4	Reserved.																				
3	<p>cHTVicDistMode: cHT VicBlk and VicBlkClean packet distribution mode. Read-write. 1=Route cHT VicBlk and VicBlkClean packets that are destined to DstNode across the link specified by DstLnk; DstLnk must specify only 1 link and must be ganged (sublink 0); If the coherent link pair distribution widget is enabled then DstLnk must be the master link; can be enabled when coherent link pair distribution is enabled or disabled; cHTReqDistEn=1 enables VicBlk and VicBlkClean distribution and is recommended to be 1; cHTRspDistEn=1 enables TgtDone and SrcDone for VicBlk and VicBlkClean and is recommended to be 1; cHTPrbDistEn must not be enabled. See 2.9.3.2.4.3 [Victim Distribution Mode].</p>																				
2	<p>cHTPrbDistEn: coherent link probe distribution enable. Read-write. 1=Enable coherent link traffic distribution for the probe virtual channel; the probes affected by this bit are limited to those sourced from the local node (as opposed to being forwarded from another node). See cHTVicDistMode.</p>																				
1	<p>cHTRspDistEn: coherent link response distribution enable. Read-write. 1=Enable coherent link traffic distribution for the response virtual channel.; the responses affected by this bit are limited to responses to the request types listed for cHTReqDistEn and which are sourced from the local node (as opposed to being forwarded from another node). See cHTVicDistMode.</p>																				
0	<p>cHTReqDistEn: coherent link request distribution enable. Read-write. 1=Enable coherent link traffic distribution for the request virtual channel; the requests affected by this bit are limited to cache block transactions and directed probes which are sourced from the local node (as opposed to being forwarded from another node). See cHTVicDistMode.</p>																				

D18F0x168 Extended Link Transaction Control

Reset: 0000_0000h.

Bits	Description
31:11	Reserved.
10	DisNcHtCmdThrottle: disable IO link command throttling. Read-write. 0=The node limits generation of the first DWORD of link-defined commands to no more than one every four DWORDs of link bandwidth. If, for example, a 2-DWORD command is transmitted by the node, and there is no data that follows, then the node sends at least 2 DWORDs of NOPs (possibly including buffer release credits) before generating the next command packet. This bit applies to both Gen1 and Gen3 frequencies and protocols. This bit does not affect coherent links. Some IO devices may require this bit to be clear. 1=The node does not limit the rate at which commands are generated on IO links.
9:0	Reserved.

D18F0x16C Link Global Extended Control**Same-for-all.** Further information about these bits can be found in the Gen3 link specification.

Bits	Description																				
31:23	Reserved.																				
22:17	FullT0Time: full T0 time. Read-write. Reset: 3Ah. BIOS: 33h. Specifies the amount of time to spend in training 0 following a warm reset, frequency change, or when the full T0 training period is invoked due to expiration of the idle timer as described in D18F0x16C[ForceFullT0] . Encodings are the same as T0Time. BIOS should set FullT0Time according to the maximum T0 training time requirement for the link's far-side receiver phase recovery time as determined by characterization.																				
16	ImmUpdate: immediate update. Cold reset: 0. Read-write. Many of the link phy registers, accessed through D18F4x1[98,90,88,80] , control electrical parameters that are unsafe to change while the link is operational; so the updates to these registers are normally withheld until the link is disconnected. However, under some (testing and characterization) circumstances, it is preferable to allow these changes to occur immediately, while the link is operational. ImmUpdate provides this option. 0=Writes to most of the link phy registers do not take effect in the link phy until the next LDTSTOP_L or warm reset disconnect. Reads from a link phy register after a write return the current value not the value pending until the link disconnects and reconnects. 1=Writes to the link phy registers are passed to the phy immediately. Reads always returns the value from the most recent write.																				
15:13	ForceFullT0: force full T0 training time. Read-write. Cold reset: 000b. BIOS: IF (D18F0x[18C:170][LS2En] is set for any link) THEN 111b ELSE 000b ENDIF. This specifies the period of time that the link can be disconnected for an LDTSTOP_L before the full T0 training period is invoked. The time is measured approximately from the assertion of LDTSTOP_L until training 0 is about to start after LDTSTOP_L deassertion. If less than the specified time has expired, then training 0 specified by T0Time is used. <table border="1" data-bbox="267 1669 1047 1848"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Disabled.</td> <td>100b</td> <td>3.2 ms</td> </tr> <tr> <td>001b</td> <td>400 us</td> <td>101b</td> <td>6.4 ms</td> </tr> <tr> <td>010b</td> <td>800 us</td> <td>110b</td> <td>12.8 ms</td> </tr> <tr> <td>011b</td> <td>1.6 ms</td> <td>111b</td> <td>25.6 ms</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	Disabled.	100b	3.2 ms	001b	400 us	101b	6.4 ms	010b	800 us	110b	12.8 ms	011b	1.6 ms	111b	25.6 ms
Bits	Description	Bits	Description																		
000b	Disabled.	100b	3.2 ms																		
001b	400 us	101b	6.4 ms																		
010b	800 us	110b	12.8 ms																		
011b	1.6 ms	111b	25.6 ms																		
12:10	Reserved.																				

9	RXCalEn: receiver calibration enable. Read-write. Cold reset: 0. 1=Enable receiver offset calibration during all training 1 periods.										
8	ConnDly: connect delay. Read-write. Cold reset: 0. 1=Changes to D18F0x[E4,C4,A4,84][TransOff, EndOfChain] take effect on the next LDTSTOP_L or warm reset. For dual-node processors, BIOS should program this according to section 2.12.1.5 [Link Mapping Between Package and Node] .										
7:6	<p>InLnSt: inactive lane state. Read-write; reset-applied. Cold reset: 00b. BIOS: 01b. Specifies the state of inactive lanes of ganged links at Gen3 frequencies. Updates to this bit take effect on warm reset or LDTSTOP_L. BIOS recommendation of PHY OFF related disabling unused sub-link of a ganged link; see 2.12.1.5 [Link Mapping Between Package and Node].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Same as warm reset except CAD is logical 0.</td> </tr> <tr> <td>01b</td> <td>Same as PHY OFF.</td> </tr> <tr> <td>10b</td> <td>Same as operational; CTL and CAD transmit undefined scrambled data.</td> </tr> <tr> <td>11b</td> <td>Same as disconnected per D18F0x[18C:170][LS2En].</td> </tr> </tbody> </table> <ul style="list-style-type: none"> • If (InLnSt=11b and D18F0x[18C:170][LS2En]=0): <ul style="list-style-type: none"> • A link width increase cannot be done after a frequency change, unless it is accompanied by another frequency change or a warm reset. • If InLnSt=00b, InLnSt=01b, or (InLnSt=11b and D18F0x[18C:170][LS2En]=1): <ul style="list-style-type: none"> • A width change via LDTSTOP_L incurs FullT0Time. 	Bits	Description	00b	Same as warm reset except CAD is logical 0.	01b	Same as PHY OFF.	10b	Same as operational; CTL and CAD transmit undefined scrambled data.	11b	Same as disconnected per D18F0x[18C:170][LS2En] .
Bits	Description										
00b	Same as warm reset except CAD is logical 0.										
01b	Same as PHY OFF.										
10b	Same as operational; CTL and CAD transmit undefined scrambled data.										
11b	Same as disconnected per D18F0x[18C:170][LS2En] .										
5:0	<p>T0Time: training 0 time. Read-write. Cold reset: 3Ah. Specifies the amount of time to spend in training 0 when exiting the disconnected state. See ForceFullT0, D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1][Ls2ExitTime], and 2.12.6 [Link LDTSTOP_L Disconnect-Reconnect].</p> <table border="1"> <thead> <tr> <th>T0Time[5:4]</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>(T0Time[3:0] * 0.1) us, range: 0.0 to 1.5 us</td> </tr> <tr> <td>01b</td> <td>(T0Time[3:0] * 0.5) us, range: 0.0 to 7.5 us</td> </tr> <tr> <td>10b</td> <td>(T0Time[3:0] * 2.0) us, range: 0.0 to 30 us</td> </tr> <tr> <td>11b</td> <td>(0h<=T0Time[3:0]<=Ah): (T0Time[3:0] * 20) us, range 0.0 to 200 us (Bh<=T0Time[3:0]<=Fh): Reserved</td> </tr> </tbody> </table> <p>BIOS: IF (D18F0x170[LS2En] D18F0x174[LS2En] D18F0x178[LS2En] D18F0x17C[LS2En] D18F0x180[LS2En] D18F0x184[LS2En] D18F0x188[LS2En] D18F0x18C[LS2En]) THEN 26h ELSE 14h ENDIF. BIOS should set T0Time according to the T0 training time requirement for the links' far-side receiver phase recovery time. See D18F0x[18C:170][LS2En].</p>	T0Time[5:4]	Description	00b	(T0Time[3:0] * 0.1) us, range: 0.0 to 1.5 us	01b	(T0Time[3:0] * 0.5) us, range: 0.0 to 7.5 us	10b	(T0Time[3:0] * 2.0) us, range: 0.0 to 30 us	11b	(0h<=T0Time[3:0]<=Ah): (T0Time[3:0] * 20) us, range 0.0 to 200 us (Bh<=T0Time[3:0]<=Fh): Reserved
T0Time[5:4]	Description										
00b	(T0Time[3:0] * 0.1) us, range: 0.0 to 1.5 us										
01b	(T0Time[3:0] * 0.5) us, range: 0.0 to 7.5 us										
10b	(T0Time[3:0] * 2.0) us, range: 0.0 to 30 us										
11b	(0h<=T0Time[3:0]<=Ah): (T0Time[3:0] * 20) us, range 0.0 to 200 us (Bh<=T0Time[3:0]<=Fh): Reserved										

D18F0x[18C:170] Link Extended Control

These registers provide control for each link. They are mapped to the links as follows:

Table 146: [Register Mapping for D18F0x\[18C:170\]](#)

Register	Function
D18F0x170	Link 0 Sublink 0
D18F0x174	Link 1 Sublink 0
D18F0x178	Link 2 Sublink 0
D18F0x17C	Link 3 Sublink 0
D18F0x180	Link 0 Sublink 1
D18F0x184	Link 1 Sublink 1
D18F0x188	Link 2 Sublink 1
D18F0x18C	Link 3 Sublink 1

Visibility of these sublink 1 registers is as specified through Ganged (bit 0) of the sublink 0 registers. If a link is ganged, only the register for sublink 0 of that link is visible and it applies to the whole link. Further information about these bits can be found in the Gen3 link specification.

Bits	Description																																				
31:14	Reserved.																																				
13:12	<p>LaneSel: lanes select. Read-write. Cold reset: 00b. This field only exists in the sublink 0 registers; in the sublink 1 registers, these bits are reserved. For unganged links, they apply to both sublinks. This field specifies how receive (RX) lanes are translated into transmit (TX) lanes for links that are in ILM. The translation varies with link width. Given the RX order specified below, the TX order varies with LaneSel as follows:</p> <table border="0"> <tr> <td><u>Bits</u></td> <td><u>16-bit link</u></td> <td><u>8-bit link</u></td> </tr> <tr> <td></td> <td>RX={CTL1, CAD[15:8], CTL0, CAD[7:0]}</td> <td>RX={CTL0, CAD[7:0]}</td> </tr> <tr> <td>00b</td> <td>Same as RX.</td> <td>Same as RX</td> </tr> <tr> <td>01b</td> <td>TX={CAD[12:8], CTL0, CAD[7:0], CTL1, CAD[15:13]}</td> <td>TX={CAD[6:0], CTL0, CAD[7]}</td> </tr> <tr> <td>10b</td> <td>TX={CTL0, CAD[7:0], CTL1, CAD[15:8]}</td> <td>TX={CAD[4:0], CTL0, CAD[7:5]}</td> </tr> <tr> <td>11b</td> <td>TX={CAD[4:0], CTL1, CAD[15:8], CTL0, CAD[7:5]}</td> <td>TX={CAD[2:0], CTL0, CAD[7:3]}</td> </tr> </table> <table border="0"> <tr> <td><u>Bits</u></td> <td><u>4-bit link</u></td> <td><u>2-bit link</u></td> </tr> <tr> <td></td> <td>RX = {CTL0, CAD[3:0]}</td> <td>RX = {CTL0, CAD[1:0]}</td> </tr> <tr> <td>00b</td> <td>Same as RX.</td> <td>Same as RX.</td> </tr> <tr> <td>01b</td> <td>TX={CAD[3:0], CTL0}</td> <td>TX={CAD[1:0], CTL0}</td> </tr> <tr> <td>10b</td> <td>TX={CAD[2:0], CTL0, CAD[3]}</td> <td>TX={CAD[0], CTL0, CAD[1]}</td> </tr> <tr> <td>11b</td> <td>TX={CAD[1:0], CTL0, CAD[3:2]}</td> <td>Reserved</td> </tr> </table> <p>Values 01b and 11b are not useful at Gen1 frequencies because the link cannot be trained unless the CTL lanes line up.</p> <p>In BIST mode on 16-bit links, LaneSel[1] selects which sublink is received by the BIST engine. 0=sublink 0, 1=sublink 1; LaneSel[1:0] also causes the receive path of the BIST engine to reverse the translation for 8-bit or smaller links.</p>	<u>Bits</u>	<u>16-bit link</u>	<u>8-bit link</u>		RX={CTL1, CAD[15:8], CTL0, CAD[7:0]}	RX={CTL0, CAD[7:0]}	00b	Same as RX.	Same as RX	01b	TX={CAD[12:8], CTL0, CAD[7:0], CTL1, CAD[15:13]}	TX={CAD[6:0], CTL0, CAD[7]}	10b	TX={CTL0, CAD[7:0], CTL1, CAD[15:8]}	TX={CAD[4:0], CTL0, CAD[7:5]}	11b	TX={CAD[4:0], CTL1, CAD[15:8], CTL0, CAD[7:5]}	TX={CAD[2:0], CTL0, CAD[7:3]}	<u>Bits</u>	<u>4-bit link</u>	<u>2-bit link</u>		RX = {CTL0, CAD[3:0]}	RX = {CTL0, CAD[1:0]}	00b	Same as RX.	Same as RX.	01b	TX={CAD[3:0], CTL0}	TX={CAD[1:0], CTL0}	10b	TX={CAD[2:0], CTL0, CAD[3]}	TX={CAD[0], CTL0, CAD[1]}	11b	TX={CAD[1:0], CTL0, CAD[3:2]}	Reserved
<u>Bits</u>	<u>16-bit link</u>	<u>8-bit link</u>																																			
	RX={CTL1, CAD[15:8], CTL0, CAD[7:0]}	RX={CTL0, CAD[7:0]}																																			
00b	Same as RX.	Same as RX																																			
01b	TX={CAD[12:8], CTL0, CAD[7:0], CTL1, CAD[15:13]}	TX={CAD[6:0], CTL0, CAD[7]}																																			
10b	TX={CTL0, CAD[7:0], CTL1, CAD[15:8]}	TX={CAD[4:0], CTL0, CAD[7:5]}																																			
11b	TX={CAD[4:0], CTL1, CAD[15:8], CTL0, CAD[7:5]}	TX={CAD[2:0], CTL0, CAD[7:3]}																																			
<u>Bits</u>	<u>4-bit link</u>	<u>2-bit link</u>																																			
	RX = {CTL0, CAD[3:0]}	RX = {CTL0, CAD[1:0]}																																			
00b	Same as RX.	Same as RX.																																			
01b	TX={CAD[3:0], CTL0}	TX={CAD[1:0], CTL0}																																			
10b	TX={CAD[2:0], CTL0, CAD[3]}	TX={CAD[0], CTL0, CAD[1]}																																			
11b	TX={CAD[1:0], CTL0, CAD[3:2]}	Reserved																																			
11	<p>ILMEN: internal loopback mode (ILM) enable. Read-write. Cold reset: 0. 1=ILM enabled on the next LDTSTOP Disconnect or warm reset. Cleared by hardware upon the subsequent LDTSTOP Disconnect or warm reset. D18F4x1[9C,94,8C,84]_x[DF,CF][*XmtRdPtr, *RcvRdPtr] must be 0 (the default) when ILM mode is used.</p>																																				
10	<p>BistEn: built-in self test (BIST) enable. Read-write. Cold reset: 0. 1=The link BIST engine is enabled on the next LDTSTOP Disconnect or warm reset. Cleared by hardware upon the subsequent LDTSTOP Disconnect or warm reset.</p>																																				
9	Reserved.																																				
8	<p>LS2En: LDTSTOP mode 2 enable. Read-write. Cold reset: 0. BIOS: 1. 0=Use LS1 mode for power reduction when the link is disconnected. 1=Use LS2 mode.</p>																																				
7:4	Reserved.																																				

3	ScrambleEn: scrambling enable. Read-write. Cold reset: 0. 1=Scrambling enable. Updates to this bit take effect on warm reset and LDTSTOP. Software must clear this bit when transitioning from Gen3 to Gen1 protocol. See Table 117 .
2:1	Reserved.
0	Ganged. Read-write; read-only 1 if the bit corresponding to the link in D18F3xE8[UnGangEn] is 0. 0=The link is unganged; this register is visible for both sublinks. 1=The link is ganged; only the sublink 0 register is visible. This value is initialized after a cold reset, based on the ganging state determined by hardware (see 2.12.1.1 [Ganging And Unganging]). Writes to this bit take effect on the next warm reset; reads reflect the last value written (rather than the current state of the link). This bit only exists in the sublink 0 registers. BIOS: see 2.12.1.5 [Link Mapping Between Package and Node] .

D18F0x1A0 Link Initialization Status

Bits	Description																																				
31	InitStatusValid: initialization status valid. Read-only. 1=Indicates that the rest of the information in this register is valid for all links; each link is either not connected or the initialization is complete.																																				
30:24	Reserved.																																				
23:16	IntLnkRoute: internal link routing. Read-only. Value: Product-specific. Defines an 8 bit vector, with one bit per sublink, indicating whether the links are connected to internal nodes or external nodes. Valid for all links regardless of the initialization status if D18F3xE8[MultiNodeCpu]=1 . 1=Internal node. 0=External node. See D18F0x[E8,C8,A8,88][Freq] and D18F4x1[9C,94,8C,84]_x[D5,C5] . See Figure 15: [Dual-Node Processor Link Diagram] . <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>link 0 sublink 0</td> <td>[4]</td> <td>link 0 sublink 1</td> </tr> <tr> <td>[1]</td> <td>link 1 sublink 0</td> <td>[5]</td> <td>link 1 sublink 1</td> </tr> <tr> <td>[2]</td> <td>link 2 sublink 0</td> <td>[6]</td> <td>link 2 sublink 1</td> </tr> <tr> <td>[3]</td> <td>link 3 sublink 0</td> <td>[7]</td> <td>link 3 sublink 1</td> </tr> </tbody> </table>	Bit	Description	Bit	Description	[0]	link 0 sublink 0	[4]	link 0 sublink 1	[1]	link 1 sublink 0	[5]	link 1 sublink 1	[2]	link 2 sublink 0	[6]	link 2 sublink 1	[3]	link 3 sublink 0	[7]	link 3 sublink 1																
Bit	Description	Bit	Description																																		
[0]	link 0 sublink 0	[4]	link 0 sublink 1																																		
[1]	link 1 sublink 0	[5]	link 1 sublink 1																																		
[2]	link 2 sublink 0	[6]	link 2 sublink 1																																		
[3]	link 3 sublink 0	[7]	link 3 sublink 1																																		
15:0	NcAndInitComplete. Read-only. These bits provide duplicate versions of D18F0x[F8,D8,B8,98][Nc, InitComplete] . <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>link 0 sublink 0 InitComplete</td> <td>[8]</td> <td>link 0 sublink 1 InitComplete</td> </tr> <tr> <td>[1]</td> <td>link 0 sublink 0 NC</td> <td>[9]</td> <td>link 0 sublink 1 NC</td> </tr> <tr> <td>[2]</td> <td>link 1 sublink 0 InitComplete</td> <td>[10]</td> <td>link 1 sublink 1 InitComplete</td> </tr> <tr> <td>[3]</td> <td>link 1 sublink 0 NC</td> <td>[11]</td> <td>link 1 sublink 1 NC.</td> </tr> <tr> <td>[4]</td> <td>link 2 sublink 0 InitComplete</td> <td>[12]</td> <td>link 2 sublink 1 InitComplete</td> </tr> <tr> <td>[5]</td> <td>link 2 sublink 0 NC</td> <td>[13]</td> <td>link 2 sublink 1 NC.</td> </tr> <tr> <td>[6]</td> <td>link 3 sublink 0 InitComplete</td> <td>[14]</td> <td>link 3 sublink 1 InitComplete</td> </tr> <tr> <td>[7]</td> <td>link 3 sublink 0 NC</td> <td>[15]</td> <td>link 3 sublink 1 NC</td> </tr> </tbody> </table>	Bit	Description	Bit	Description	[0]	link 0 sublink 0 InitComplete	[8]	link 0 sublink 1 InitComplete	[1]	link 0 sublink 0 NC	[9]	link 0 sublink 1 NC	[2]	link 1 sublink 0 InitComplete	[10]	link 1 sublink 1 InitComplete	[3]	link 1 sublink 0 NC	[11]	link 1 sublink 1 NC.	[4]	link 2 sublink 0 InitComplete	[12]	link 2 sublink 1 InitComplete	[5]	link 2 sublink 0 NC	[13]	link 2 sublink 1 NC.	[6]	link 3 sublink 0 InitComplete	[14]	link 3 sublink 1 InitComplete	[7]	link 3 sublink 0 NC	[15]	link 3 sublink 1 NC
Bit	Description	Bit	Description																																		
[0]	link 0 sublink 0 InitComplete	[8]	link 0 sublink 1 InitComplete																																		
[1]	link 0 sublink 0 NC	[9]	link 0 sublink 1 NC																																		
[2]	link 1 sublink 0 InitComplete	[10]	link 1 sublink 1 InitComplete																																		
[3]	link 1 sublink 0 NC	[11]	link 1 sublink 1 NC.																																		
[4]	link 2 sublink 0 InitComplete	[12]	link 2 sublink 1 InitComplete																																		
[5]	link 2 sublink 0 NC	[13]	link 2 sublink 1 NC.																																		
[6]	link 3 sublink 0 InitComplete	[14]	link 3 sublink 1 InitComplete																																		
[7]	link 3 sublink 0 NC	[15]	link 3 sublink 1 NC																																		

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.

D18F0x1E0 Coherent Link Pair Traffic Distribution

Reset: 0000_0000h. See 2.9.3.2.4 [Link Traffic Distribution] for details about link traffic distribution.

Bits	Description																				
31:29	AltSel3: alternate select 3. See: AltSel0.																				
28:26	MasterSel3: master select 3. See: MasterSel0.																				
25	Asym3: asymmetric 3. See: Asym0.																				
24	DistEn3: distribution enable 3. See: DistEn0.																				
23:21	AltSel2: alternate select 2. See: AltSel0.																				
20:18	MasterSel2: master select 2. See: MasterSel0.																				
17	Asym2: asymmetric 2. See: Asym0.																				
16	DistEn2: distribution enable 2. See: DistEn0.																				
15:13	AltSel1: alternate select 1. See: AltSel0.																				
12:10	MasterSel1: master select 1. See: MasterSel0.																				
9	Asym1: asymmetric 1. See: Asym0.																				
8	DistEn1: distribution enable 1. See: DistEn0.																				
7:5	AltSel0: alternate select 0. Select for alternate link of pair 0. See: MasterSel0.																				
4:2	MasterSel0: master select 0. Read-write. Select for master link of pair 0. <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>link 0 sublink 0</td> <td>100b</td> <td>link 0 sublink 1</td> </tr> <tr> <td>001b</td> <td>link 1 sublink 0</td> <td>101b</td> <td>link 1 sublink 1</td> </tr> <tr> <td>010b</td> <td>link 2 sublink 0</td> <td>110b</td> <td>link 2 sublink 1</td> </tr> <tr> <td>011b</td> <td>link 3 sublink 0</td> <td>111b</td> <td>link 3 sublink 1</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	link 0 sublink 0	100b	link 0 sublink 1	001b	link 1 sublink 0	101b	link 1 sublink 1	010b	link 2 sublink 0	110b	link 2 sublink 1	011b	link 3 sublink 0	111b	link 3 sublink 1
Bits	Description	Bits	Description																		
000b	link 0 sublink 0	100b	link 0 sublink 1																		
001b	link 1 sublink 0	101b	link 1 sublink 1																		
010b	link 2 sublink 0	110b	link 2 sublink 1																		
011b	link 3 sublink 0	111b	link 3 sublink 1																		
1	Asym0: asymmetric 0. Read-write. 1=Link pair 0 is asymmetric.																				
0	DistEn0: distribution enable 0. Read-write. 1=Enables traffic distribution for link pair 0.																				

3.4 Device [1F:18]h Function 1 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.8 [Configuration Space].

D18F1x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1601h.
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[17C:140,7C:40] DRAM Base/Limit

The following 8 sets of registers specify the destination node of DRAM address ranges:

Table 147: Register Mapping for D18F1x[17C:140,7C:40]

Function	Base Low	Limit Low	Base High	Limit High
Range 0	D18F1x40	D18F1x44	D18F1x140	D18F1x144
Range 1	D18F1x48	D18F1x4C	D18F1x148	D18F1x14C
Range 2	D18F1x50	D18F1x54	D18F1x150	D18F1x154
Range 3	D18F1x58	D18F1x5C	D18F1x158	D18F1x15C
Range 4	D18F1x60	D18F1x64	D18F1x160	D18F1x164
Range 5	D18F1x68	D18F1x6C	D18F1x168	D18F1x16C
Range 6	D18F1x70	D18F1x74	D18F1x170	D18F1x174
Range 7	D18F1x78	D18F1x7C	D18F1x178	D18F1x17C

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 [2.9.3 \[NB Routing\]](#).

DRAM mapping rules:

- Transaction addresses are within the defined range if:
{DramBase[47:24], 00_0000h} <= address[47:0] <= {DramLimit[47:24], FF_FFFh}.
- 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 [2.9.3.1.1 \[DRAM and MMIO Memory Space\]](#).

Hoisting. When memory hoisting is enabled in a node via [D18F1xF0](#)[DramHoleValid], the corresponding BaseAddr/LimitAddr should be configured to account for the memory hoisted above the hole. See [2.10.7 \[Memory Hoisting\]](#).

Node interleave. DRAM may be mapped as continuous regions for each node or it may be interleaved between nodes. See [2.10.6.3 \[Node Interleaving\]](#).

D18F1x[78,70,68,60,58,50,48,40] DRAM Base Low

IF ([D18F2x118](#)[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0000_0000h.

Table 148: [Register Mapping](#) for [D18F1x\[78,70,68,60,58,50,48,40\]](#)

Register	Function
D18F1x40	Range 0
D18F1x48	Range 1
D18F1x50	Range 2
D18F1x58	Range 3
D18F1x60	Range 4
D18F1x68	Range 5
D18F1x70	Range 6
D18F1x78	Range 7

Bits	Description														
31:16	DramBase[39:24]: DRAM base address register bits[39:24]. DramBase[47:24] = { D18F1x[178,170,168,160,158,150,148,140] [DramBase[47:40]], D18F1x[78,70,68,60,58,50,48,40] [DramBase[39:24]]}.														
15:11	Reserved.														
10:8	IntlvEn[2:0]: interleave enable. Enables interleaving on a 4-KB boundary between memory on different nodes. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>No interleave</td> </tr> <tr> <td>001b</td> <td>Interleave on A[12] (2 nodes)</td> </tr> <tr> <td>010b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>Interleave on A[12] and A[13] (4 nodes)</td> </tr> <tr> <td>110b-100b</td> <td>Reserved</td> </tr> <tr> <td>111b</td> <td>Interleave on A[12], A[13], and A[14] (8 nodes)</td> </tr> </tbody> </table> The value of this field is required to match D18F1x124 [DramIntlvEn].	Bits	Description	000b	No interleave	001b	Interleave on A[12] (2 nodes)	010b	Reserved	011b	Interleave on A[12] and A[13] (4 nodes)	110b-100b	Reserved	111b	Interleave on A[12], A[13], and A[14] (8 nodes)
Bits	Description														
000b	No interleave														
001b	Interleave on A[12] (2 nodes)														
010b	Reserved														
011b	Interleave on A[12] and A[13] (4 nodes)														
110b-100b	Reserved														
111b	Interleave on A[12], A[13], and A[14] (8 nodes)														
7:2	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.														

D18F1x[178,170,168,160,158,150,148,140] DRAM Base High

Table 149: Register Mapping for D18F1x[178,170,168,160,158,150,148,140]

Register	Function
D18F1x140	Range 0
D18F1x148	Range 1
D18F1x150	Range 2
D18F1x158	Range 3
D18F1x160	Range 4
D18F1x168	Range 5
D18F1x170	Range 6
D18F1x178	Range 7

Bits	Description
31:8	Reserved.
7:0	DramBase[47:40]: DRAM base address register bits[47:40]. See: D18F1x[78,70,68,60,58,50,48,40][DramBase[39:24]] .

D18F1x[7C,74,6C,64,5C,54,4C,44] DRAM Limit Low

IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF.

Table 150: Register Mapping for D18F1x[7C,74,6C,64,5C,54,4C,44]

Register	Function
D18F1x44	Range 0
D18F1x4C	Range 1
D18F1x54	Range 2
D18F1x5C	Range 3
D18F1x64	Range 4
D18F1x6C	Range 5
D18F1x74	Range 6
D18F1x7C	Range 7

Bits	Description
31:16	DramLimit[39:24]: DRAM limit address register bits[39:24]. Reset: FFFFh. DramLimit[47:24] = {D18F1x[17C,174,16C,164,15C,154,14C,144][DramLimit[47:40]], D18F1x[7C,74,6C,64,5C,54,4C,44][DramLimit[39:24]]} .
15:11	Reserved.
10:8	IntlvSel: interleave select. Reset: 000b. Specifies the values of address bits A[14:12] to use with the Interleave Enable field (IntlvEn[2:0]) to determine which 4-KB blocks are routed to this region. IntlvSel[0] corresponds to A[12]; IntlvSel[1] corresponds to A[13]; IntlvSel[2] corresponds to A[14].

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.

D18F1x[17C,174,16C,164,15C,154,14C,144] DRAM Limit High

IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF.

Table 151: Register Mapping for D18F1x[17C,174,16C,164,15C,154,14C,144]

Register	Function
D18F1x144	Range 0
D18F1x14C	Range 1
D18F1x154	Range 2
D18F1x15C	Range 3
D18F1x164	Range 4
D18F1x16C	Range 5
D18F1x174	Range 6
D18F1x17C	Range 7

Bits	Description
31:8	Reserved.
7:0	DramLimit[47:40]: DRAM limit address register bits[47:40]. Reset: 00h. See D18F1x[7C,74,6C,64,5C,54,4C,44][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 mapping from memory addresses to the corresponding node and IO link for MMIO transactions. Address ranges are specified by 12 sets of base/limit registers.

Table 152: Register Mapping for D18F1x[1CC:180,BC:80]

Function	MMIO Base Low	MMIO Limit Low	MMIO Base/Limit High
Range 0	D18F1x80	D18F1x84	D18F1x180
Range 1	D18F1x88	D18F1x8C	D18F1x184
Range 2	D18F1x90	D18F1x94	D18F1x188
Range 3	D18F1x98	D18F1x9C	D18F1x18C
Range 4	D18F1xA0	D18F1xA4	D18F1x190
Range 5	D18F1xA8	D18F1xAC	D18F1x194
Range 6	D18F1xB0	D18F1xB4	D18F1x198
Range 7	D18F1xB8	D18F1xBC	D18F1x19C
Range 8	D18F1x1A0	D18F1x1A4	D18F1x1C0

Table 152: Register Mapping for D18F1x[1CC:180,BC:80]

Range 9	D18F1x1A8	D18F1x1AC	D18F1x1C4
Range 10	D18F1x1B0	D18F1x1B4	D18F1x1C8
Range 11	D18F1x1B8	D18F1x1BC	D18F1x1CC

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 2.9.3 [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[17C:140,7C: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 2.9.3.1.1 [DRAM and MMIO Memory Space].

D18F1x[1B8,1B0,1A8,1A0,B8,B0,A8,A0,98,90,88,80] MMIO Base Low

Table 153: 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
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. $\text{MMIOBase}[47:16] = \{\text{D18F1x}[1\text{CC}:1\text{C}0,19\text{C}:180][\text{MMIOBase}[47:40]], \text{MMIOBase}[39:16]\}$.

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 154: 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][DsNpReqLmt] 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 <i>HyperTransport™ IO Link Specification</i> 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	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.										
5:4	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. <table border="1"> <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 MMIO range, this field specifies the destination node ID.										

D18F1x[1CC:1C0,19C:180] MMIO Base/Limit High

Table 155: Register Mapping for 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
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 2.9.3 [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 2.9.3.1.2 [IO Space].

D18F1x[D8,D0,C8,C0] IO-Space Base

Table 156: 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 ever 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 157: 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 ungang, 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 ganged, 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 2.9.3 [NB Routing].

Table 158: 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 2.8 [Configuration Space] in the case of MMIO accesses) are within the defined range if:

({BusNumBase[7:0]} <= BusNo <= {BusNumLimit[7:0]}) & (DevCmpEn==0); or
 ({BusNumBase[4:0]} <= Device <= {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 2.9.3.1.3 [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.										
15:11	Reserved.										
10	DstSubLink: destination sublink . Read-write. Reset: 0. When a link is unanged, this bit specifies the destination sublink of the link specified by D18F1x[EC:E0][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.										
9:8	DstLink: destination link ID . Read-write. Reset: 0. For transactions within the this configuration-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										
7	Reserved.										
6:4	DstNode: destination node ID bits . Read-write. Reset: 0. For transactions within the this configuration-space range, this field specifies the destination node ID.										
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). This is used to enable multiple IO links to be configured as Bus 0.										
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. Reset: 0000_0000h. See 2.10.7 [Memory Hoisting].

Bits	Description
31:24	DramHoleBase[31:24]: DRAM hole base address . Specifies the base address of the IO hole, below the 4GB address level, that is used in memory hoisting. Normally, DramHoleBase >= MSRC001_001A[TOM[31:24]] .
23:16	Reserved.
15:7	DramHoleOffset[31:23]: DRAM hole offset address . When DramHoleValid =1, this offset is subtracted from the physical address of certain accesses in forming the normalized address.
6:2	Reserved.

1	DramMemHoistValid: dram memory hoist valid. 1=Memory hoisting for the address range is enabled in one or more nodes of the coherent fabric. 0=Memory hoisting is not enabled for any node. This bit should be set in all nodes of the coherent fabric if DramHoleValid=1 in any node.
0	DramHoleValid: dram hole valid. 1=Memory hoisting is enabled in the node. 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. If node interleaving is employed, then this should be set in all nodes. 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 ungangd, 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 gangd, 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"> <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. Specifies the set of DCT NB P-state 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>1xb</td> <td>Reserved</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	1xb	Reserved
Bits	Description								
00b	NB P-state 0								
01b	NB P-state 1								
1xb	Reserved								
3:1	Unused.								
0	<p>DctCfgSel: DRAM controller configuration select. Read-write. Specifies DCT controller to which accesses are routed. 0=DCT 0. 1=DCT 1. See 2.10.1 [DCT Configuration Registers].</p>								

D18F1x120 DRAM Base System Address

IF ([D18F2x118\[LockDramCfg\]](#)) THEN Read-only. ELSE Read-write. ENDIF. 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:

$\{\text{DramBaseAddr}[47:27], 000_0000\text{h}\} \leq \text{Addr}[47:0] \leq \{\text{DramLimitAddr}[47:27], 7\text{FF_FFFFh}\};$

DRAM accesses to the local node that are outside of this range are master aborted. This range is also used to specify the range of DRAM covered by the scrubber (see [D18F3x58](#) and [D18F3x5C](#)).

A base/limit pair cannot include any part of the reserved address range between `FD_0000_0000h` to `FF_FFFF_FFFFh`.

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`.

If node interleaving is invoked, then `DramBaseAddr` should be zero in all the nodes and `DramLimitAddr` should be the top of memory in all nodes. Based on the value of `DramIntlvEn`, the normalized address to the DCTs is modified to remove the affected address bits between `A[17:12]`; e.g., if 8-node interleave is invoked, then `DramIntlvEn` is set to 111b and the normalized address to the DCTs removes `A[14:12]` to become $\{\text{A}[47:15], \text{A}[11:0]\}$. See [2.10.6.3 \[Node Interleaving\]](#).

Bits	Description
31:24	Reserved.
23:21	DramIntlvSel: interleave select. Specifies the values of address bits <code>A[14:12]</code> that are routed to the local node when node interleaving is enabled. <code>IntlvSel[0]</code> corresponds to <code>A[12]</code> ; <code>IntlvSel[1]</code> corresponds to <code>A[13]</code> ; <code>IntlvSel[2]</code> corresponds to <code>A[14]</code> .
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														
31:24	Reserved														
23:21	DramIntlvEn[2:0] . Reset: 0. Specifies interleaving on a 4-KB boundary between DRAM on different nodes. The bits are encoded as follows: <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>No interleave</td> </tr> <tr> <td>001b</td> <td>Interleave on A[12] (2 nodes)</td> </tr> <tr> <td>010b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>Interleave on A[12] and A[13] (4 nodes)</td> </tr> <tr> <td>110b-100b</td> <td>Reserved</td> </tr> <tr> <td>111b</td> <td>Interleave on A[12], A[13], and A[14] (8 nodes)</td> </tr> </tbody> </table> The value of this field is required to match D18F1x[17C:140,7C:40][IntlvEn] .	Bits	Description	000b	No interleave	001b	Interleave on A[12] (2 nodes)	010b	Reserved	011b	Interleave on A[12] and A[13] (4 nodes)	110b-100b	Reserved	111b	Interleave on A[12], A[13], and A[14] (8 nodes)
Bits	Description														
000b	No interleave														
001b	Interleave on A[12] (2 nodes)														
010b	Reserved														
011b	Interleave on A[12] and A[13] (4 nodes)														
110b-100b	Reserved														
111b	Interleave on A[12], A[13], and A[14] (8 nodes)														
20:0	DramLimitAddr[47:27] . Reset: 1F_FFFFh.														

3.5 Device [1F:18]h Function 2 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.8 [Configuration Space].

D18F2x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID . Read-only. Value: 1602h.
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 2.10.1 [DCT Configuration Registers].

Table 159: Register Mapping for D18F2x[5C:40]_dct[1:0]

Register	Function
D18F2x40_dct[1:0]	0
D18F2x44_dct[1:0]	1
D18F2x48_dct[1:0]	2
D18F2x4C_dct[1:0]	3
D18F2x50_dct[1:0]	4
D18F2x54_dct[1:0]	5
D18F2x58_dct[1:0]	6
D18F2x5C_dct[1:0]	7

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 2.10 [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. These are associated with DIMM numbers, and CKE signals as follows:

Table 160: DIMM, Chip Select, and Register Mapping

Base Address Registers	Mask Register	Logical DIMM ¹		Chip Select Pad ²	CKE Pad ²
		N	R4		
F2x40	F2x60	0	0	MEMCS[1:0]_L[0]	MEMCKE[1:0][0]
F2x44				MEMCS[1:0]_L[1]	MEMCKE[1:0][1]
F2x48	F2x64	1	1	MEMCS[1:0]_L[2]	MEMCKE[1:0][0]
F2x4C				MEMCS[1:0]_L[3]	MEMCKE[1:0][1]
F2x50	F2x68	2	0	MEMCS[1:0]_L[4]	MEMCKE[1:0][0]
F2x54				MEMCS[1:0]_L[5]	MEMCKE[1:0][1]
F2x58	F2x6C	3	1	MEMCS[1:0]_L[6]	MEMCKE[1:0][0]
F2x5C				MEMCS[1:0]_L[7]	MEMCKE[1:0][1]

1. N=Normal.
R4=Four-rank registered DIMM. See D18F2x94_dct[1:0][FourRankRDimm1, FourRankRDimm0].

2. See 2.10.2 [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 2.9 [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 [2.10.6.1 \[Chip Select Interleaving\]](#). 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[21:13] & ~AddrMask[i][21:13]) } ==
  { (BaseAddr[i][38:27] & ~AddrMask[i][38:27]),
    (BaseAddr[i][21:13] & ~AddrMask[i][21:13]) } );
```

Bits	Description
31	Reserved.
30:19	BaseAddr[38:27]: normalized physical base address bits [38:27].
18:14	Reserved.
13:5	BaseAddr[21:13]: normalized physical base address bits [21:13].
4	Reserved.
3	OnDimmMirror: on-DIMM mirroring (ODM) enabled. BIOS: RankMap . 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 D18F2x7C_dct[1:0][SendMrsCmd] when D18F2x7C_dct[1:0][EndDramInit] = 0. This bit is expected to be set for the odd numbered rank of UDIMMs and LRDIMMs if SPD byte 63 indicates that address mapping is mirrored. See also 2.10.5.7.1.1 [DDR3 MR Initialization] , 2.10.5.7.1.2 [Software Control Word Initialization] , and F0RC14[AddressMirror] . For LRDIMMs, BIOS should not enable this mode in the DCT until after 2.10.5.7.1 [Software DDR3 Device Initialization] is complete. 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 CSEnable=1 and TestFail=1 as mutually exclusive.
1	Spare: spare rank. This bit identifies the chip select associated with the spare rank. See 2.10.9 [On-Line Spare] .
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 [2.10.1](#)

[DCT Configuration Registers]. See D18F2x[5C:40]_dct[1:0].

Table 161: Register Mapping for D18F2x[6C:60]_dct[1:0]

Register	Function
D18F2x60_dct[1:0]	0
D18F2x64_dct[1:0]	1
D18F2x68_dct[1:0]	2
D18F2x6C_dct[1:0]	3

Table 162: BIOS Recommendations for D18F2x[6C:60]_dct[1:0][RankDef]

Condition				D18F2x[6C:60]_dct[1:0]
DIMM	NumRanks	DramCapacity	NumDimmSlots	RankDef
SODIMM UDIMM RDIMM	-	-	-	00b
LRDIMM	000b	-	-	01b
	001b	-	-	01b
	011b	-	1, 2	01b
	011b	-	3	10b
	100b	0100b	-	11b
	100b	-	1, 2	10b
	100b	-	3	11b

Bits	Description										
31	Reserved.										
30:19	AddrMask[38:27]: normalized physical address mask bits [38:27].										
18:14	Reserved.										
13:5	AddrMask[21:13]: normalized physical address mask bits [21:13].										
4:2	Reserved.										
1:0	RankDef: rank definition. Specifies the rank definition. <table border="0" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>SODIMM, UDIMM, or RDIMM</td> </tr> <tr> <td>01b</td> <td>LRDIMM 1x rank multiply</td> </tr> <tr> <td>10b</td> <td>LRDIMM 2x rank multiply</td> </tr> <tr> <td>11b</td> <td>LRDIMM 4x rank multiply</td> </tr> </tbody> </table>	Bits	Description	00b	SODIMM, UDIMM, or RDIMM	01b	LRDIMM 1x rank multiply	10b	LRDIMM 2x rank multiply	11b	LRDIMM 4x rank multiply
Bits	Description										
00b	SODIMM, UDIMM, or RDIMM										
01b	LRDIMM 1x rank multiply										
10b	LRDIMM 2x rank multiply										
11b	LRDIMM 4x rank multiply										

D18F2x78_dct[1:0] DRAM Control

See 2.10.1 [DCT Configuration Registers].

Bits	Description
31:18	Reserved.
17	AddrCmdTriEn: address command tristate enable. Read-write. Reset: 0. BIOS: See 2.10.5.6.1=Tristate the address, command, and bank buses when a Deselect command is issued.

16:15	Reserved.
14:0	Reserved.

D18F2x7C_dct[1:0] DRAM Initialization

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.7.1 [Software DDR3 Device Initialization].

BIOS can directly control the DRAM initialization sequence using this register. To do so, 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 initialization. BIOS should not assert LDTSTOP_L while EnDramInit is set. 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 D18F2x98_dct[1:0][DctAccessDone] = 1 before programming AssertCke=1 and DeassertMemRstX=1. BIOS must clear this bit after DRAM initialization is complete. BIOS must not set this bit on a DCT with no attached DIMMs. See 2.10.5.7.1 [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 D18F2xA8_dct[1:0][CtrlWordCS] . This bit is cleared by hardware after the command completes. This bit is valid only when D18F2x90_dct[1:0][UnbuffDimm] = 0. Reserved if D18F2x78_dct[1:0][AddrCmdTriEn] =1
29	SendZQCmd: send ZQ command. Read; write-1-only; cleared-by-hardware. 1=The DCT sends the ZQCL 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. Reserved if D18F2x7C_dct[1:0][EnDramInit] = 0 or D18F2x78_dct[1:0][AddrCmdTriEn] =1 .
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 MrsAddress and MrsBank fields of this register. This bit is cleared by hardware after the command completes. See 2.10.5.7.1.1 [DDR3 MR Initialization]. Reserved if D18F2x78_dct[1:0][AddrCmdTriEn] =1 .
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 (~EnDrAmInit ~D18F2x90_dct[1:0][UnbuffDimm]); otherwise MRS commands are sent to all chip selects. <table border="1"> <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 2.10.1 [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 width information of the DIMM. Table 163 and Table 164 show the bit numbers for each position.

Bits	Description
31:16	Reserved.
15:12	Dimm3AddrMap: DIMM 3 address map.
11:8	Dimm2AddrMap: DIMM 2 address map.
7:4	Dimm1AddrMap: DIMM 1 address map.
3:0	Dimm0AddrMap: DIMM 0 address map.

Table 163: 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	1GB	512Mb, x4	16	15	14	Row	x	x	x	17	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
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

Table 163: 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	
1000b	2GB	1Gb, x4	16	15	14	Row	x	x	17	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
1001b	4GB	2Gb, x4	16	15	14	Row	x	17	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
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	4Gb, x4 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	16GB	8Gb, x4	17	16	15	Row	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18
						Col	x	x	14	x	13	AP	12	11	10	9	8	7	6	5	4	3

Table 164: DDR3 LRDIMM Rank Multiplication Address Mapping

Bits	CS Size ¹	Device size, width	Row Address			
			17	16	15	14
0011b - 0000b		Reserved				
0101b	2GB, 4GB	1Gb, x8	x	x	31	30
0110b		Reserved				
0111b	4GB, 8GB	2Gb, x8	x	32	31	V ²
1000b	4GB, 8GB	1Gb, x4	x	x	32	31
1001b	8GB, 16GB	2Gb, x4	x	33	32	V ²
1010b	8GB, 16GB	4Gb, x8	33	32	V ²	V ²
1011b	16GB, 32GB	4Gb, x4	34	33	V ²	V ²
1100b		Reserved				

1. 2x, 4x rank multiplication. See 2.10.11.1 [LRDIMM Rank Multiplication].
 2. V = Valid. See Table 163.

D18F2x84_dct[1:0] DRAM MRS

Reset: 0000_0004h. See 2.10.1 [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_dct[1:0][PPD]. See 2.10.5.7 [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</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][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	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																										
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: 00b. Specifies the number of sequential beats of DQ related to one read or write command.</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>11b-01b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	8 beats	11b-01b	Reserved																						
<u>Bits</u>	<u>Description</u>																												
00b	8 beats																												
11b-01b	Reserved																												

D18F2x88_dct[1:0] DRAM Timing Low

See [2.10.1 \[DCT Configuration Registers\]](#).

Bits	Description														
31:30	Reserved.														
29:24	<p>MemClkDis: MEMCLK disable. Read-write. Reset: 3Fh. BIOS: Table 103-Table 109. 1=Disable the MEMCLK. 0=Enable MEMCLK. All enabled clocks should be 0; all disabled and N/A clocks should be 1. See 2.10.2 [DDR Pad to Processor Pin Mapping].</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 [2.10.1 \[DCT Configuration Registers\]](#).

Bits	Description
31:19	Reserved.

18	DisAutoRefresh: disable automatic refresh. Read-write. BIOS: See 2.10.5.6 . 1=Automatic refresh is disabled.										
17:16	Tref: refresh rate. Read-write. BIOS: See 2.10.5.4 and 2.10.10 . 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 [2.10.1 \[DCT Configuration Registers\]](#).

Bits	Description										
31:28	Reserved.										
27	DisDllShutdownSR: disable DLL shutdown in self-refresh mode. Read-write; <i>Same-for-all</i> . Reset: 1. BIOS: IF (AM3r2) THEN 1 ELSE See 2.10.5.3.2.1 . ENDIF. 1=Disable the power saving feature of shutting down DDR phy DLLs during DRAM self refresh. 0=Shutdown DLLs during DRAM self refresh.										
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: 1. 1=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] . This field ignored for LRDIMM; stagger refresh selected automatically by LRDIMM.										
23	ForceAutoPchg: force auto precharging. Read-write. Reset: 0. See 2.10.5.6 . 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 2.10.5.6 . 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	DimmEccEn: DIMM ECC enable. Read-write. Reset: 0. 1=ECC checking is capable of being enabled for all DIMMs on the DRAM controller (through D18F3x44[DrumEccEn]). This bit should not be set unless all populated DIMMs support ECC check bits. 0=ECC checking is disabled on the DRAM controller.										

18	PendRefPayback: pending refresh payoff. 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 2.10.5.8 [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 LRDIMMs.										
15:12	X4Dimm: x4 (by 4) DIMMs. Read-write. Reset: 0. BIOS: Set the corresponding DIMM bit if (<code>DeviceWidth == 000b && ~LRDIMM</code>). Each of these bits specifies whether the corresponding DIMM (as defined by <code>D18F2x[5C:40]_dct[1:0]</code> [DRAM CS Base Address]) is a x4 DIMM or not. The DRAM controller requires this information to make decisions about DIMM signaling. 1=x4 DIMM present. 0=x4 DIMM not present. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>DIMM 0</td> </tr> <tr> <td>[1]</td> <td>DIMM 1</td> </tr> <tr> <td>[2]</td> <td>DIMM 2</td> </tr> <tr> <td>[3]</td> <td>DIMM 3</td> </tr> </tbody> </table>	Bit	Description	[0]	DIMM 0	[1]	DIMM 1	[2]	DIMM 2	[3]	DIMM 3
Bit	Description										
[0]	DIMM 0										
[1]	DIMM 1										
[2]	DIMM 2										
[3]	DIMM 3										
11:9	Reserved.										
8	ParEn: parity enable. Read-write. Reset: 0. BIOS: See 2.10.5.7.1. 1=Enables address parity computation output, PAR, and enables the parity error input, ERR. This bit is valid only when UnbuffDimm==0. See <code>D18F2xA8_dct[1:0]</code> [ExtendedParityEn].										
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 2.10.5.8 [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 <code>MR0[PPD]=1</code> if Fast exit precharge powerdown mode is desired. See <code>D18F2x94_dct[1:0][1:0]</code> [PowerDownEn] and <code>D18F2x84_dct[1:0]</code> [PchgPDMModeSel].										
0	Reserved.										

D18F2x94_dct[1:0] DRAM Configuration High

See 2.10.1 [DCT Configuration Registers].

Table 165: Valid Values for D18F2x94_dct[1:0][MemClkFreq]

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)
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)
1Fh-17h	Reserved

Bits	Description										
31:28	Reserved.										
27:24	<p>DcqBypassMax: DRAM controller queue bypass maximum. Read-write. Reset: 0h. BIOS: 2.10.5.6. 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.</p> <table border="1"> <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>1h</td> <td>The oldest request may be bypassed no more than 1 time.</td> </tr> <tr> <td>Eh-2h</td> <td>The oldest request may be bypassed no more than <DcqBypassMax> time.</td> </tr> <tr> <td>Fh</td> <td>The oldest request may be bypassed no more than 15 times.</td> </tr> </tbody> </table>	Bits	Description	0h	No bypass; the oldest request is never bypassed.	1h	The oldest request may be bypassed no more than 1 time.	Eh-2h	The oldest request may be bypassed no more than <DcqBypassMax> time.	Fh	The oldest request may be bypassed no more than 15 times.
Bits	Description										
0h	No bypass; the oldest request is never bypassed.										
1h	The oldest request may be bypassed no more than 1 time.										
Eh-2h	The oldest request may be bypassed no more than <DcqBypassMax> time.										
Fh	The oldest request may be bypassed no more than 15 times.										
23	<p>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][ProcOdt] for ODT definitions. Changes to this bit must be performed prior to DRAM initialization.</p>										

22	<p>BankSwizzleMode: bank swizzle mode. Read-write. Reset: 0. BIOS: See 2.10.5.6. 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 163 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 163 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. BIOS: Table 73-Table 79. 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.</p>
19	<p>DcqArbBypassEn: DRAM controller arbiter bypass enable. Read-write. Reset: 1. 0=DCQ entries are always passed through the arbiter. 1=Bypass the arbitration logic when there is only one entry in the DRAM controller queue entry.</p>
18	<p>FourRankRDimm0: four rank registered DIMM 0. Read-write. Reset: 0. BIOS: 0. 1=Four-rank registered DIMM or four-rank LRDIMM is connected to the channel CS 0, 1, 4, 5. See Table 160 [DIMM, Chip Select, and Register Mapping]. Changes the parameter selection for D18F2x214_dct[1:0][TwrwrSdDc, TwrwrDd] and D18F2x218_dct[1:0][TrdrdSdDc, TrdrdDd].</p>
17	<p>FourRankRDimm1: four rank registered DIMM 1. Read-write. Reset: 0. BIOS: 0. 1=Four-rank registered DIMM or four-rank LRDIMM is connected to the channel CS 2, 3, 6, 7. See Table 160 [DIMM, Chip Select, and Register Mapping]. Changes the parameter selection for D18F2x214_dct[1:0][TwrwrSdDc, TwrwrDd] and D18F2x218_dct[1:0][TrdrdSdDc, TrdrdDd].</p>
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 160 [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	PowerDownEn: power down mode enable. Read-write. Reset: 0. BIOS: See 2.10.5.6. 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] , and D18F2x248_dct[1:0][PchgPDEnDelay] .										
14	DisDramInterface: disable the DRAM interface. Read-write. Reset: 0. BIOS: See 2.10.5. 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. If this bit is set, D18F2x90_dct[1:0][DimmEccEn] and D18F2x90_dct[1:0][ParEn] should not be set to avoid spurious MCA errors.										
13	Reserved.										
12	Reserved.										
11:10	ZqcsInterval: ZQ calibration short interval. Read-write. Reset: 00b. BIOS: See 2.10.5.6. This field specifies the programmable interval for the controller to send out the DRAM ZQ calibration short command. <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	MemClkFreqVal: memory clock frequency valid. Read-write. Reset: 0. BIOS: See 2.10.5.3.2. 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. 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.										
6:5	Reserved.										
4:0	MemClkFreq: memory clock frequency. Read-write. Reset: 000b. Specifies the frequency and rate of the DRAM interface (MEMCLK). See: Table 165 [Valid Values for D18F2x94_dct[1:0][MemClkFreq]] . The rate is twice the frequency. See D18F5x84[DdrMaxRate] and D18F5x84[DdrMaxRateEnf] .										

D18F2x98_dct[1:0] DRAM Controller Additional Data Offset

Reset: 8000_0000h. See 2.10.1 [\[DCT Configuration Registers\]](#).

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:
 - Write the register number to [D18F2x98_dct\[1:0\]\[DctOffset\]](#) with [D18F2x98_dct\[1:0\]\[DctAccessWrite\]=0](#).
 - Poll until [D18F2x98_dct\[1:0\]\[DctAccessDone\]=1](#).
 - Read the register contents from [D18F2x9C_dct\[1:0\]](#).
- Writes:
 - Write all 32 bits of register data to [D18F2x9C_dct\[1:0\]](#) (individual byte writes are not supported).
 - Write the register number to [D18F2x98_dct\[1:0\]\[DctOffset\]](#) with [D18F2x98_dct\[1:0\]\[DctAccessWrite\]=1](#).

Write]=1.

- Poll until `D18F2x98_dct[1:0][DctAccessDone]=1`. This ensures that the contents of the write have been delivered to the phy.

Writes to any register in this additional address space causes the FIFO pointers to be reset.

Bits	Description
31	DctAccessDone: DRAM controller access complete. Read-only. 1=The access to the register is complete. 0=The access is still in progress.
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 2.10.1 [DCT Configuration Registers].

D18F2x9C_x0000_0000_dct[1:0] DRAM Output Driver Compensation Control

BIOS: [Table 73-Table 79](#).

Bits	Description												
31	Reserved.												
30:28	ProcOdt: processor on-die termination. Read-write. Cold reset: 000b. Specifies the resistance of the on-die termination resistors. This field is valid only when <code>D18F2x94_dct[1:0][ProcOdtDis]=0</code> . <table border="1" data-bbox="267 1102 698 1312"> <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												
27:23	Reserved.												

22:20	<p>DqsDrvStren: DQS drive strength. Read-write. Cold reset: 011b. Specifies the drive strength of the DQS pins.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0.75x</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> <p>The DM[8:0] and DQS[17:9] functions share pins on the DIMM connector. The function selection is applied based on whether the DIMM is populated with by-4 (x4) DRAM devices, in which case the DQS[17:9] function is applied, or not (x8 or x16 DRAM devices), in which case the DM[8:0] function is applied. However, the DM function is associated with the data pin group and should therefore be controlled DataDrvStren. While the processor supports concurrent population of x4 and non-x4 DIMMs, the determination as to which field controls the drive strength of these pins is applied statically based on these rules:</p> <ul style="list-style-type: none"> • If all DIMMs of a channel are populated with non-x4 devices, DataDrvStren is applied. • If any DIMMs of a channel are populated with x4 devices, DqsDrvStren is applied. 	<u>Bits</u>	<u>Description</u>	000b	0.75x	001b	1.0x	010b	1.25x	011b	1.5x	111b-100b	Reserved
<u>Bits</u>	<u>Description</u>												
000b	0.75x												
001b	1.0x												
010b	1.25x												
011b	1.5x												
111b-100b	Reserved												
19	Reserved.												
18:16	<p>DataDrvStren: data drive strength. Read-write. Cold reset: 011b. This field specifies the drive strength of the DRAM data pins.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>0.75x</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> <p>See the note in DqsDrvStren regarding how this field may be applied to DM signals as well.</p>	<u>Bits</u>	<u>Description</u>	000b	0.75x	001b	1.0x	010b	1.25x	011b	1.5x	111b-100b	Reserved
<u>Bits</u>	<u>Description</u>												
000b	0.75x												
001b	1.0x												
010b	1.25x												
011b	1.5x												
111b-100b	Reserved												
15	Reserved.												
14:12	<p>ClkDrvStren: MEMCLK drive strength. Read-write. Cold reset: 011b. This field specifies the drive strength of the MEMCLK pins.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	000b	1.0x	001b	1.25x	010b	1.5x	011b	2.0x	111b-100b	Reserved
<u>Bits</u>	<u>Description</u>												
000b	1.0x												
001b	1.25x												
010b	1.5x												
011b	2.0x												
111b-100b	Reserved												
11	Reserved.												
10:8	<p>AddrCmdDrvStren: address/command drive strength. Read-write. Cold reset: 011b. This field specifies the drive strength of the address, RAS, CAS, WE, bank and parity pins.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	000b	1.0x	001b	1.25x	010b	1.5x	011b	2.0x	111b-100b	Reserved
<u>Bits</u>	<u>Description</u>												
000b	1.0x												
001b	1.25x												
010b	1.5x												
011b	2.0x												
111b-100b	Reserved												
7	Reserved.												

6:4	CsOdtDrvStren: CS/ODT drive strength. Read-write. Cold reset: 011b. This field specifies the drive strength of the CS and ODT 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												
3	Reserved.												
2:0	CkeDrvStren: CKE drive strength. Read-write. Cold reset: 011b. This field specifies the drive strength of the CKE 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_x0000_0[3:0]0[3:1]_dct[1:0] DRAM Write Data Timing

BIOS: See 2.10.5.8.4.

Table 166: Register Mapping for D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0]

Register	Function
D18F2x9C_x0000_0001_dct[1:0]	DIMM 0 Bytes 3-0
D18F2x9C_x0000_0002_dct[1:0]	DIMM 0 Bytes 7-4
D18F2x9C_x0000_0003_dct[1:0]	DIMM 0 ECC
D18F2x9C_x0000_0101_dct[1:0]	DIMM 1 Bytes 3-0
D18F2x9C_x0000_0102_dct[1:0]	DIMM 1 Bytes 7-4
D18F2x9C_x0000_0103_dct[1:0]	DIMM 1 ECC
D18F2x9C_x0000_0201_dct[1:0]	DIMM 2 Bytes 3-0
D18F2x9C_x0000_0202_dct[1:0]	DIMM 2 Bytes 7-4
D18F2x9C_x0000_0203_dct[1:0]	DIMM 2 ECC
D18F2x9C_x0000_0301_dct[1:0]	DIMM 3 Bytes 3-0
D18F2x9C_x0000_0302_dct[1:0]	DIMM 3 Bytes 7-4
D18F2x9C_x0000_0303_dct[1:0]	DIMM 3 ECC

Table 167: Field Mapping for D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0]

Register	Bits			
	31:24	23:16	15:8	7:0
D18F2x9C_x0000_0[3:0]01_dct[1:0]	Byte3	Byte2	Byte1	Byte0
D18F2x9C_x0000_0[3:0]02_dct[1:0]	Byte7	Byte6	Byte5	Byte4
D18F2x9C_x0000_0[3:0]03_dct[1:0]	Reserved	Reserved	Reserved	ECC

These registers control the timing of write DQ with respect to MEMCLK and allow transmit DQS to be cen-

tered in the data eye. The delay starts 1 UI before the rising edge of MEMCLK corresponding to the CAS-write-latency. See 2.10.5.8 [DRAM Training]. WrDatGrossDly must be programmed for a given DIMM and lane such that WrDatDly - WrDqsDly <= 0.5 MEMCLKs.

Bits	Description										
31:29	WrDatGrossDly: write data gross delay. See: D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0][7:5].										
28:24	WrDatFineDly: write data fine delay. See: D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0][4:0].										
23:21	WrDatGrossDly: write data gross delay. See: D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0][7:5].										
20:16	WrDatFineDly: write data fine delay. See: D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0][4:0].										
15:13	WrDatGrossDly: write data gross delay. See: D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0][7:5].										
12:8	WrDatFineDly: write data fine delay. See: D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0][4:0].										
7:5	<p>WrDatGrossDly: write data gross delay. Read-write. Reset: 0.</p> <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	<p>WrDatFineDly: write data fine delay. Read-write. Cold reset: 0.</p> <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] DRAM Address/Command Timing Control

BIOS: Table 73-Table 79.

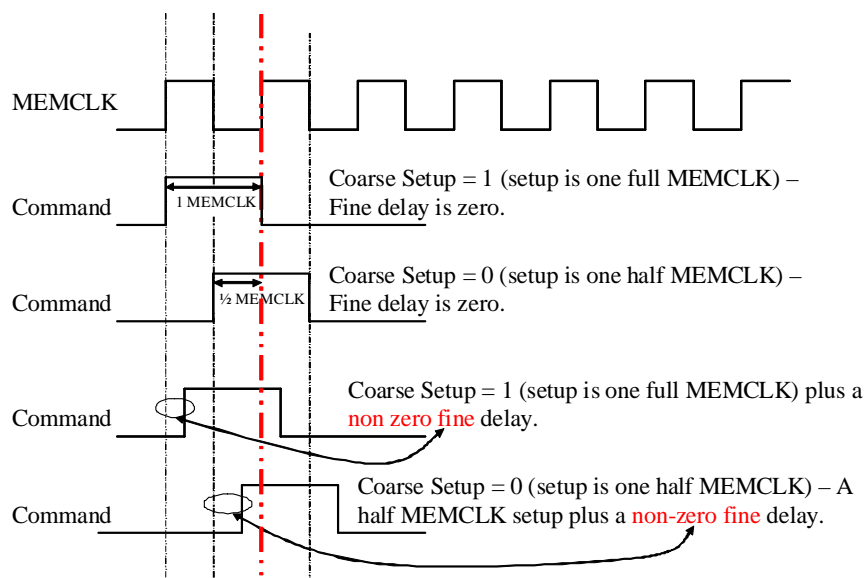


Figure 18: Address/Command Timing at the Processor Pins

This register controls the timing of the address, command, chip select, ODT and clock enable pins with respect to MEMCLK as shown in Figure 18. See 2.10.5.3.2 [DRAM Channel Frequency Change] and 2.10.5.3.3 [Phy Fence Programming]. 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. 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: address/command fine delay. Specifies the time that the address and command pins are delayed from the default setup time. See: CkeFineDelay.								
15:14	Reserved.								
13	CsOdtSetup: CS/ODT setup time. Selects the default setup time for the CS and ODT pins versus MEMCLK. See: CkeSetup.								
12:8	CsOdtFineDelay: CS/ODT fine delay. Specifies the time that the CS and ODT pins are delayed from the default setup time. See: CkeFineDelay.								
7:6	Reserved.								
5	CkeSetup: CKE setup time. Read-write. Reset: 0. 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. Specifies the time that the CKE pins are delayed from the default setup time. <table border="1" style="margin-left: 20px;"> <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><CkeFineDelay>/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	<CkeFineDelay>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay
Bits	Description								
00h	0/64 MEMCLK delay								
1Eh-01h	<CkeFineDelay>/64 MEMCLK delay								
1Fh	31/64 MEMCLK delay								

D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0] DRAM Read DQS Timing

BIOS: See 2.10.5.8.4.

Table 168: Register Mapping for D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0]

Register	Function
D18F2x9C_x0000_0005_dct[1:0]	DIMM 0 Bytes 3-0
D18F2x9C_x0000_0006_dct[1:0]	DIMM 0 Bytes 7-4
D18F2x9C_x0000_0007_dct[1:0]	DIMM 0 ECC
D18F2x9C_x0000_0105_dct[1:0]	DIMM 1 Bytes 3-0
D18F2x9C_x0000_0106_dct[1:0]	DIMM 1 Bytes 7-4
D18F2x9C_x0000_0107_dct[1:0]	DIMM 1 ECC
D18F2x9C_x0000_0205_dct[1:0]	DIMM 2 Bytes 3-0
D18F2x9C_x0000_0206_dct[1:0]	DIMM 2 Bytes 7-4
D18F2x9C_x0000_0207_dct[1:0]	DIMM 2 ECC

Table 168: Register Mapping for D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0]

D18F2x9C_x0000_0305_dct[1:0]	DIMM 3 Bytes 3-0
D18F2x9C_x0000_0306_dct[1:0]	DIMM 3 Bytes 7-4
D18F2x9C_x0000_0307_dct[1:0]	DIMM 3 ECC

Table 169: Field Mapping for D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0]

Register	Bits			
	29:25	21:17	13:9	5:1
D18F2x9C_x0000_0[3:0]05_dct[1:0]	Byte3	Byte2	Byte1	Byte0
D18F2x9C_x0000_0[3:0]06_dct[1:0]	Byte7	Byte6	Byte5	Byte4
D18F2x9C_x0000_0[3:0]07_dct[1:0]	Reserved	Reserved	Reserved	ECC

These registers control the timing of read (input) DQS signals with respect to DQ.

Bits	Description								
31:30	Reserved.								
29:25	RdDqsTime: read DQS timing control. See: D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0][5:1].								
24:22	Reserved.								
21:17	RdDqsTime: read DQS timing control. See: D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0][5:1].								
16:14	Reserved.								
13:9	RdDqsTime: read DQS timing control. See: D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0][5:1].								
8:6	Reserved.								
5:1	RdDqsTime: read DQS timing control. Read-write. Cold reset: 0Fh. <table border="1" style="margin-left: 20px;"> <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><RdDqsTime>/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	<RdDqsTime>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay
Bits	Description								
00h	0/64 MEMCLK delay								
1Eh-01h	<RdDqsTime>/64 MEMCLK delay								
1Fh	31/64 MEMCLK delay								
0	Reserved.								

D18F2x9C_x0000_0008_dct[1:0] DRAM Phy Control

Cold reset: 0208_0000h.

Bits	Description
31:14	Reserved.
13	DqsRcvTrEn: DQS receiver training enable. Read-write. BIOS: See 2.10.5.8.2. 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. BIOS: See 2.10.5.8.1. 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 2.10.5.5.5 [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. BIOS: See 2.10.5.3.3. Specifies the flop to be used for phy based fence training. See PhyFenceTrEn. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</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>	Bits	Description	00b	PRE flop	01b	RxDll flop	10b	TxDll flop	11b	TxPad flop
Bits	Description										
00b	PRE flop										
01b	RxDll flop										
10b	TxDll flop										
11b	TxPad flop										
5:4	TrDimmSel: training DIMM select. Read-write. BIOS: See 2.10.5.8.1, 2.10.5.8.2. Specifies which DIMM is to be trained. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>DIMM 0</td> </tr> <tr> <td>01b</td> <td>DIMM 1</td> </tr> <tr> <td>10b</td> <td>DIMM 2</td> </tr> <tr> <td>11b</td> <td>DIMM 3</td> </tr> </tbody> </table>	Bits	Description	00b	DIMM 0	01b	DIMM 1	10b	DIMM 2	11b	DIMM 3
Bits	Description										
00b	DIMM 0										
01b	DIMM 1										
10b	DIMM 2										
11b	DIMM 3										
3	PhyFenceTrEn: phy fence training enable. Read-write. BIOS: See 2.10.5.3.3. 1=Initiate phy based fence training. 0=Stop the phy based fence training engine.										
2	TrNibbleSel: training nibble select. Read-write. BIOS: See 2.10.5.8.1, 2.10.5.8.2. Specifies nibbles of each DIMM data and ECC byte trained during write levelization and DQS receiver enable training. 0=Lower nibbles. 1=Upper nibbles.										
1	Reserved.										
0	WrtLvTrEn: write levelization training enable. Read-write. BIOS: See 2.10.5.8.1. 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

BIOS: See 2.10.5.3.

Bits	Description
31	DynModeChange: dynamic mode change. RAZ; write. Reset: 0. 1=Phy enters the state specified by PhySelfRefreshMode.
30:24	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 fence settings and signal tri-state functionality. See 2.10.2 [DDR Pad to Processor Pin Mapping].

Bits	Description
31	Reserved.
30:26	FenceThresholdTxDll: phy fence threshold transmit DLL. Read-write. Cold reset: 13h. BIOS: See 2.10.5.3.3. 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 2.10.5.3.3. 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 2.10.5.3.3. 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																		
15:14	Reserved.																		
13:12	CKETri: CKE tri-state. Read-write. Cold reset: 00b. BIOS: Table 103-Table 109. 0=The CKE signals are not tri-stated. 1=Tri-state CKE signals from the processor. <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> </tbody> </table>	Bit	Pad	[0]	MEMCKE[1,0][0]	[1]	MEMCKE[1,0][1]												
Bit	Pad																		
[0]	MEMCKE[1,0][0]																		
[1]	MEMCKE[1,0][1]																		
11:8	ODTTri: ODT tri-state. Read-write. Cold reset: 0h. BIOS: Table 103-Table 109. 0=The ODT signals are not tri-stated unless directed to by the DCT. 1=Tri-state ODT signals from the processor. <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. BIOS: Table 103-Table 109. 0=The chip select signals are not tri-stated unless directed to by the DCT. 1=Tri-state chip selects signals from the processor. <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] 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. BIOS: See 2.10.5.10. Specifies the number of PCLKs that the DLL standby signal must deassert prior to a DLL relock event or before read traffic is sent to the receive DLLs.
23	Reserved.
22:20	RxCPUpdPeriod: receive charge pump period. Read-write. BIOS: See 2.10.5.10. Specifies the number of DLL relocks required to keep the receive DLLs locked for the period where there is no read traffic.
19:16	RxMaxDurDllNoLock: receive maximum duration DLL no lock. Read-write. BIOS: See 2.10.5.6. Specifies the number of PCLK cycles that occur before the phy DLLs relock. A DLL relock occurs every $2^{\text{RxMaxDurDllNoLock}}$ if there are no reads during the period. 0=DLL power saving disabled.
15:10	Reserved.
9:8	TxDLLWakeupTime: transmit DLL wakeup time. Read-write. BIOS: See 2.10.5.10. 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.
7	Reserved.
6:4	TxCPUpdPeriod: transmit charge pump DLL wakeup time. Read-write. BIOS: See 2.10.5.10. Specifies the number of DLL relocks required to keep the TxDLLs locked for the period where there is no write traffic.
3:0	TxMaxDurDllNoLock: transmit maximum duration DLL no lock. Read-write. BIOS: See 2.10.5.6. Specifies the number of PCLK cycles that occur before the phy DLLs relock. A DLL relock occurs every $2^{\text{TxMaxDurDllNoLock}}$ if there are no writes during the period. 0=DLL power saving disabled.

D18F2x9C_x0000_00[2A:10]_dct[1:0] DRAM DQS Receiver Enable Timing

BIOS: See 2.10.5.8.2.

Table 170: Register Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]

Register	Function
D18F2x9C_x0000_0010_dct[1:0]	DIMM 0 Bytes 1-0
D18F2x9C_x0000_0011_dct[1:0]	DIMM 0 Bytes 3-2
D18F2x9C_x0000_0012_dct[1:0]	DIMM 0 ECC
D18F2x9C_x0000_0013_dct[1:0]	DIMM 1 Bytes 1-0
D18F2x9C_x0000_0014_dct[1:0]	DIMM 1 Bytes 3-2
D18F2x9C_x0000_0015_dct[1:0]	DIMM 1 ECC
D18F2x9C_x0000_0016_dct[1:0]	DIMM 2 Bytes 1-0
D18F2x9C_x0000_0017_dct[1:0]	DIMM 2 Bytes 3-2
D18F2x9C_x0000_0018_dct[1:0]	DIMM 2 ECC
D18F2x9C_x0000_0019_dct[1:0]	DIMM 3 Bytes 1-0
D18F2x9C_x0000_001A_dct[1:0]	DIMM 3 Bytes 3-2
D18F2x9C_x0000_001B_dct[1:0]	DIMM 3 ECC
D18F2x9C_x0000_0020_dct[1:0]	DIMM 0 Bytes 5-4
D18F2x9C_x0000_0021_dct[1:0]	DIMM 0 Bytes 7-6
D18F2x9C_x0000_0023_dct[1:0]	DIMM 1 Bytes 5-4

Table 170: Register Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]

D18F2x9C_x0000_0024_dct[1:0]	DIMM 1 Bytes 7-6
D18F2x9C_x0000_0026_dct[1:0]	DIMM 2 Bytes 5-4
D18F2x9C_x0000_0027_dct[1:0]	DIMM 2 Bytes 7-6
D18F2x9C_x0000_0029_dct[1:0]	DIMM 3 Bytes 5-4
D18F2x9C_x0000_002A_dct[1:0]	DIMM 3 Bytes 7-6

Table 171: Field Mapping for D18F2x9C_x0000_00[2A:10]_dct[1:0]

Register	Bits	
	25:16	9:0
D18F2x9C_x0000_001[9,6,3,0]_dct[1:0]	Byte1	Byte0
D18F2x9C_x0000_001[A,7,4,1]_dct[1:0]	Byte3	Byte2
D18F2x9C_x0000_001[B,8,5,2]_dct[1:0]	Reserved	ECC
D18F2x9C_x0000_002[9,6,3,0]_dct[1:0]	Byte5	Byte4
D18F2x9C_x0000_002[A,7,4,1]_dct[1:0]	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 2.10.5.8 [DRAM Training]. These delay registers must be programmed such that across all DIMMs and lanes $MAX(DqsRcvEnDelay) - MIN(DqsRcvEnDelay) \leq 7 UI$.

Bits	Description								
31:26	Reserved.								
25:21	DqsRcvEnGrossDelay: DQS receiver enable gross delay. See: D18F2x9C_x0000_00[2A:10]_dct[1:0][9:5].								
20:16	DqsRcvEnFineDelay: DQS receiver enable fine delay. See: D18F2x9C_x0000_00[2A:10]_dct[1:0][4:0].								
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>0/2 MEMCLK delay</td> </tr> <tr> <td>1Eh-01h</td> <td><DqsRcvEnGrossDelay>/2 MEMCLK delay</td> </tr> <tr> <td>1Fh</td> <td>31/2 MEMCLK delay</td> </tr> </tbody> </table>	Bits	Description	00h	0/2 MEMCLK delay	1Eh-01h	<DqsRcvEnGrossDelay>/2 MEMCLK delay	1Fh	31/2 MEMCLK delay
Bits	Description								
00h	0/2 MEMCLK delay								
1Eh-01h	<DqsRcvEnGrossDelay>/2 MEMCLK delay								
1Fh	31/2 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] DRAM DQS Write Timing

BIOS: See 2.10.5.8.1. Each of these registers control the DQS timing delay for write commands relative to MEMCLK. The delay starts at the rise edge of MEMCLK corresponding to the CAS-writelateness. Each control includes a gross timing field and a fine timing field, the sum of which is the total delay.

Table 172: Register Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]

Register	Function
D18F2x9C_x0000_0030_dct[1:0]	DIMM 0 Bytes 1-0
D18F2x9C_x0000_0031_dct[1:0]	DIMM 0 Bytes 3-2
D18F2x9C_x0000_0032_dct[1:0]	DIMM 0 ECC
D18F2x9C_x0000_0033_dct[1:0]	DIMM 1 Bytes 1-0
D18F2x9C_x0000_0034_dct[1:0]	DIMM 1 Bytes 3-2
D18F2x9C_x0000_0035_dct[1:0]	DIMM 1 ECC
D18F2x9C_x0000_0036_dct[1:0]	DIMM 2 Bytes 1-0
D18F2x9C_x0000_0037_dct[1:0]	DIMM 2 Bytes 3-2
D18F2x9C_x0000_0038_dct[1:0]	DIMM 2 ECC
D18F2x9C_x0000_0039_dct[1:0]	DIMM 3 Bytes 1-0
D18F2x9C_x0000_003A_dct[1:0]	DIMM 3 Bytes 3-2
D18F2x9C_x0000_003B_dct[1:0]	DIMM 3 ECC
D18F2x9C_x0000_0040_dct[1:0]	DIMM 0 Bytes 5-4
D18F2x9C_x0000_0041_dct[1:0]	DIMM 0 Bytes 7-6
D18F2x9C_x0000_0043_dct[1:0]	DIMM 1 Bytes 5-4
D18F2x9C_x0000_0044_dct[1:0]	DIMM 1 Bytes 7-6
D18F2x9C_x0000_0046_dct[1:0]	DIMM 2 Bytes 5-4
D18F2x9C_x0000_0047_dct[1:0]	DIMM 2 Bytes 7-6
D18F2x9C_x0000_0049_dct[1:0]	DIMM 3 Bytes 5-4
D18F2x9C_x0000_004A_dct[1:0]	DIMM 3 Bytes 7-6

Table 173: Field Mapping for D18F2x9C_x0000_00[4A:30]_dct[1:0]

Register	Bits	
	23:16	7:0
D18F2x9C_x0000_003[9,6,3,0]_dct[1:0]	Byte1	Byte0
D18F2x9C_x0000_003[A,7,4,1]_dct[1:0]	Byte3	Byte2
D18F2x9C_x0000_003[B,8,5,2]_dct[1:0]	Reserved	ECC
D18F2x9C_x0000_004[9,6,3,0]_dct[1:0]	Byte5	Byte4
D18F2x9C_x0000_004[A,7,4,1]_dct[1:0]	Byte7	Byte6

Bits	Description
31:29	Reserved.
28:24	Reserved.
23:21	WrDqsGrossDly: DQS write gross delay. See: D18F2x9C_x0000_00[4A:30]_dct[1:0][7:5] .
20:16	WrDqsFineDly: DQS write fine delay. See: D18F2x9C_x0000_00[4A:30]_dct[1:0][4:0] .
15:13	Reserved.
12:8	Reserved.

7:5	WrDqsGrossDly: DQS write 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><WrDqsGrossDly>/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	<WrDqsGrossDly>/2 MEMCLK delay	111b	3.5 MEMCLK delay
Bits	Description										
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.										
	<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><WrDqsFineDly>/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	<WrDqsFineDly>/64 MEMCLK delay	1Fh	31/64 MEMCLK delay		
Bits	Description										
00h	0/64 MEMCLK delay										
1Eh-01h	<WrDqsFineDly>/64 MEMCLK delay										
1Fh	31/64 MEMCLK delay										

D18F2x9C_x0000_00[52:50]_dct[1:0] DRAM Phase Recovery Control

BIOS: See 2.10.5.8. 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.

Table 174: Register Mapping for D18F2x9C_x0000_00[52:50]_dct[1:0]

Register	Function
D18F2x9C_x0000_0050_dct[1:0]	Bytes 3-0
D18F2x9C_x0000_0051_dct[1:0]	Bytes 7-4
D18F2x9C_x0000_0052_dct[1:0]	ECC

Table 175: Field Mapping for D18F2x9C_x0000_00[52:50]_dct[1:0]

Register	Bits			
	31:24	23:16	15:8	7:0
D18F2x9C_x0000_0050_dct[1:0]	Byte3	Byte2	Byte1	Byte0
D18F2x9C_x0000_0051_dct[1:0]	Byte7	Byte6	Byte5	Byte4
D18F2x9C_x0000_0052_dct[1:0]	Reserved	Reserved	Reserved	ECC

Bits	Description
31	Reserved.
30:29	PhRecGrossDly: phase recovery gross delay. See: D18F2x9C_x0000_00[52:50]_dct[1:0][6:5].
28:24	PhRecFineDly: phase recovery fine delay. See: D18F2x9C_x0000_00[52:50]_dct[1:0][4:0].
23	Reserved.
22:21	PhRecGrossDly: phase recovery gross delay. See: D18F2x9C_x0000_00[52:50]_dct[1:0][6:5].
20:16	PhRecFineDly: phase recovery fine delay. See: D18F2x9C_x0000_00[52:50]_dct[1:0][4:0].
15	Reserved.
14:13	PhRecGrossDly: phase recovery gross delay. See: D18F2x9C_x0000_00[52:50]_dct[1:0][6:5].
12:8	PhRecFineDly: phase recovery fine delay. See: D18F2x9C_x0000_00[52:50]_dct[1:0][4:0].
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.
<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.
<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,8:0]02_dct[1:0] Data Byte Transmit PreDriver Calibration

Cold reset: xxxx_xxxxh. BIOS: See 2.10.5.3.4.

Table 176: Register Mapping for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0]

Register	Function
D18F2x9C_x0D0F_0002_dct[1:0]	Byte 0
D18F2x9C_x0D0F_0102_dct[1:0]	Byte 1
D18F2x9C_x0D0F_0202_dct[1:0]	Byte 2
D18F2x9C_x0D0F_0302_dct[1:0]	Byte 3
D18F2x9C_x0D0F_0402_dct[1:0]	Byte 4
D18F2x9C_x0D0F_0502_dct[1:0]	Byte 5
D18F2x9C_x0D0F_0602_dct[1:0]	Byte 6
D18F2x9C_x0D0F_0702_dct[1:0]	Byte 7
D18F2x9C_x0D0F_0802_dct[1:0]	ECC
D18F2x9C_x0D0F_0F02_dct[1:0]	Broadcast

D18F2x9C_x0D0F_0802_dct[1:0] is broadcast to D18F2x9C_x0D0F_0[8:0]02_dct[1:0].

Table 177: Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]

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

Table 177: Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]

Bits	Description
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,8:0]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 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]]. 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 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]]. After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.

D18F2x9C_x0D0F_0[F,8:0]04_dct[1:0] Data Byte DM Configuration

Cold reset: 0000_0033h.

Table 178: Index Mapping for D18F2x9C_x0D0F_0[F,8:0]04_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]								
	0804h	0704h	0604h	0504h	0404h	0304h	0204h	0104h	0004h
0D0Fh	ECC	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 179: Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
0D0Fh	D18F2x9C_x0D0F_0[8:0]04

Bits	Description
31:14	Reserved.
13	TriDM: tri-state DM. Read-write. BIOS: 2.10.5.10. Read-write. Specifies tri-state control for the memory DM signal. 1=Signal is tri-stated. 0=Signal is not tri-stated.
12:0	Reserved.

D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0] Data Byte Transmit PreDriver Calibration 2

Cold reset: xxxx_xxxxh. BIOS: See 2.10.5.3.4.

Table 180: Register Mapping for D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0]

Register	Function
D18F2x9C_x0D0F_0006_dct[1:0]	Byte 0
D18F2x9C_x0D0F_0106_dct[1:0]	Byte 1
D18F2x9C_x0D0F_0206_dct[1:0]	Byte 2
D18F2x9C_x0D0F_0306_dct[1:0]	Byte 3
D18F2x9C_x0D0F_0406_dct[1:0]	Byte 4
D18F2x9C_x0D0F_0506_dct[1:0]	Byte 5
D18F2x9C_x0D0F_0606_dct[1:0]	Byte 6
D18F2x9C_x0D0F_0706_dct[1:0]	Byte 7
D18F2x9C_x0D0F_0806_dct[1:0]	ECC
D18F2x9C_x0D0F_0F06_dct[1:0]	Broadcast
D18F2x9C_x0D0F_000A_dct[1:0]	Byte 0
D18F2x9C_x0D0F_010A_dct[1:0]	Byte 1
D18F2x9C_x0D0F_020A_dct[1:0]	Byte 2
D18F2x9C_x0D0F_030A_dct[1:0]	Byte 3
D18F2x9C_x0D0F_040A_dct[1:0]	Byte 4
D18F2x9C_x0D0F_050A_dct[1:0]	Byte 5
D18F2x9C_x0D0F_060A_dct[1:0]	Byte 6
D18F2x9C_x0D0F_070A_dct[1:0]	Byte 7
D18F2x9C_x0D0F_080A_dct[1:0]	ECC
D18F2x9C_x0D0F_0F0A_dct[1:0]	Broadcast

D18F2x9C_x0D0F_0F06_dct[1:0] is broadcast to D18F2x9C_x0D0F_0[8:0]06_dct[1:0].

D18F2x9C_x0D0F_0F0A_dct[1:0] is broadcast to D18F2x9C_x0D0F_0[8:0]0A_dct[1:0].

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: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]] . After updating this value, BIOS must program D18F2x9C_x0D0F_0[F,8:0]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: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]] . After updating this value, BIOS must program D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][ValidTxAndPre]=1 for the change to take effect.

D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0] Data Byte DLL Clock Enable

Cold reset: 0000_0013h.

Table 181: Index Mapping for D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]								
	080Fh	070Fh	060Fh	050Fh	040Fh	030Fh	020Fh	010Fh	000Fh
0D0Fh	ECC	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 182: Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F0Fh
0D0Fh	D18F2x9C_x0D0F_0[8:0]0F

Bits	Description								
31:15	Reserved.								
14:12	<p>AlwaysEnDIIlClks: always enable DLL clocks. Read-write. BIOS: 2.10.5.3.3. 0=DLL clocks are turned off during periods of inactivity. 1=DLL clocks remain on during inactivity. Prior to programming AlwaysEnDIIlClks to a value other than 000b, D18F2x9C_x0000_000D_dct[1:0][RxMaxDurDIIlNoLock, TxMaxDurDIIlNoLock] must both be programmed to 0000b. The bits AlwaysEnDIIlClks[2:0] are mapped to DLLs as follows:</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</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	Description	[2]	TxDqs DLL	[1]	TxDq DLL	[0]	RxEn DLL
Bit	Description								
[2]	TxDqs DLL								
[1]	TxDq DLL								
[0]	RxEn DLL								
11:0	Reserved.								

D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0] Data Byte DLL Power Management

Cold reset: 0000_0000h.

Table 183: Index Mapping for D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]								
	0810h	0710h	0610h	0510h	0410h	0310h	0210h	0110h	0010h
0D0Fh	ECC	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 184: Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F10h
0D0Fh	D18F2x9C_x0D0F_0[8:0]10

Bits	Description
31:13	Reserved.

12	EnRxPadStandby: enable receiver pad standby. Read-write. BIOS: See 2.10.5.6. 1=Phy receiver standby mode is enabled to save power when not receiving data. 0=Phy receiver standby mode is disabled.
11:0	Reserved.

D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0] Data Byte DLL Configuration

Table 185: Index Mapping for D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]								
	0813h	0713h	0613h	0513h	0413h	0313h	0213h	0113h	0013h
0D0Fh	ECC	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 186: Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F13h
0D0Fh	D18F2x9C_x0D0F_0[8,7:0]13

Bits	Description
31:15	Reserved.
14	ProcOdtAdv: ProcOdt advance. Read-write. Cold reset: 1. BIOS: IF (SODIMM && DdrRate <= 1333) THEN 0 ELSE 1 ENDIF. 0=ProcOdt is asserted 1.5 PCLK before DqsRcvEn. 1=If preceding write, ProcOdt is asserted 2.5 PCLK before DqsRcvEn, else, ProcOdt is asserted 5.0 PCLK before DqsRcvEn.
13:8	Reserved.
7	RxDqsUDllPowerDown: receive DQS upper DLL power down. Read-write. Cold reset: 0. BIOS: See 2.10.5.10. 1=Power down the upper receiver DQS DLL.
6:2	Reserved.
1	DllDisEarlyU: DLL disable early upper. Read-write. Cold reset: 0. BIOS: See 2.10.5.10. 1=Disable upper receiver DQS DLL early timing for power savings.
0	DllDisEarlyL: DLL disable early lower. Read-write. Cold reset: 0. BIOS: See 2.10.5.10. 1=Disable lower receiver DQS DLL early timing for power savings.

D18F2x9C_x0D0F_0[F,8:0]1F_dct[1:0] Data Byte Receiver Configuration

Cold reset: 0000_2002h.

Table 187: Index Mapping for D18F2x9C_x0D0F_0[F,8:0]1F_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]								
	081Fh	071Fh	061Fh	051Fh	041Fh	031Fh	021Fh	011Fh	001Fh
0D0Fh	ECC	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 188: Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]1F_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F1Fh
0D0Fh	D18F2x9C_x0D0F_0[8,7:0]1F

Bits	Description										
31:5	Reserved.										
4:3	RxVioLvl: receiver voltage level. Read-write. BIOS: 2.10.5.3.1. Specifies the VDDIO voltage level. <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										
2:0	Reserved.										

D18F2x9C_x0D0F_0[F,8:0]30_dct[1:0] Data Byte DLL Configuration and Power Down

Table 189: Index Mapping for D18F2x9C_x0D0F_0[F,8:0]30_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]								
	0830h	0730h	0630h	0530h	0430h	0330h	0230h	0130h	0030h
0D0Fh	ECC	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 190: Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]30_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F30h
0D0Fh	D18F2x9C_x0D0F_0[8:0]30

Bits	Description
31:9	Reserved.
8	BlockRxDqsLock: block rx dqs lock. Read-write. Cold reset: 0. BIOS: 2.10.5.8.3. 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. Cold reset: 0. 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.
3:0	Reserved.

D18F2x9C_x0D0F_0[F,8:0]31_dct[1:0] Data Byte Fence2 Threshold

BIOS: 2.10.5.3.3.

Table 191: Index Mapping for D18F2x9C_x0D0F_0[F,8:0]31_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]								
	0831h	0731h	0631h	0531h	0431h	0331h	0231h	0131h	0031h
0D0Fh	ECC	Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0

Table 192: Broadcast Mapping for D18F2x9C_x0D0F_0[F,8:0]31_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	0F31h
0D0Fh	D18F2x9C_x0D0F_0[8:0]31

Bits	Description										
31:15	Reserved.										
14	Fence2EnableRxDll: phy fence2 enable receive DLL. Read-write. Cold reset: 0. 1=Enable the use of Fence2ThresholdRxDll for DQS receiver enable fence threshold. 0=Use D18F2x9C_x0000_000C_dct[1:0][FenceThresholdRxDll] .										
13:10	Fence2ThresholdRxDll: phy fence2 threshold DQS receiver enable. Read-write. Cold reset: 0. If Fence2EnableRxDll=1, this field specifies the fence delay threshold value used for DQS receiver enable. See Fence2ThresholdTxPad.										
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="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</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>	Bits	Description	0h	No delay	1h	1/64 MEMCLK delay	Eh-2h	<Fence2ThresholdTxPad> MEMCLK delay	Fh	15/64 MEMCLK delay
Bits	Description										
0h	No delay										
1h	1/64 MEMCLK delay										
Eh-2h	<Fence2ThresholdTxPad> MEMCLK delay										
Fh	15/64 MEMCLK delay										

D18F2x9C_x0D0F_2[2:0]02_dct[1:0] Clock Transmit PreDriver Calibration

Cold reset: xxxx_xxxxh. BIOS: See 2.10.5.3.4.

Table 193: Index Mapping for D18F2x9C_x0D0F_2[2:0]02_dct[1:0]

D18F2x98_dct[1:0][31:0]	Function
-------------------------	----------

Table 193: Index Mapping for D18F2x9C_x0D0F_2[2:0]02_dct[1:0]

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: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]] . 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: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]] . After updating this value, BIOS must program ValidTxAndPre=1 for the change to take effect.

D18F2x9C_x0D0F_[C,8,2][2:0]1F_dct[1:0] Receiver Configuration

Cold reset: 0000_2000h.

Table 194: Index Mapping for D18F2x9C_x0D0F_[C,8,2][2:0]1F_dct[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_201Fh	Clock 0
0D0F_211Fh	Clock 1
0D0F_221Fh	Clock 1
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 2.10.5.3.1. 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. BIOS: See 2.10.5.10.

Table 195: Index Mapping for D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]		
	2230h	2130h	2030h
0D0Fh	Clk 2	Clk 1	Clk 0

Table 196: Broadcast Mapping for D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0]

D18F2x98_dct[1:0][31:16]	D18F2x98_dct[1:0][15:0]
	2F30h
0D0Fh	D18F2x9C_x0D0F_2[2:0]30

Bits	Description
31:16	Reserved.
15:5	Reserved.
4	PwrDn: power down. Read-write. BIOS: IF (REG == D18F2x9C_x0D0F_2030_dct[1:0]) THEN (D18F2x88_dct[1:0][MemClkDis[0]] & D18F2x88_dct[1:0][MemClkDis[1]]) ELSEIF (REG == D18F2x9C_x0D0F_2130_dct[1:0]) THEN (D18F2x88_dct[1:0][MemClkDis[2]] & D18F2x88_dct[1:0][MemClkDis[3]]) ELSEIF (REG == D18F2x9C_x0D0F_2230_dct[1:0]) THEN (D18F2x88_dct[1:0][MemClkDis[4]] & D18F2x88_dct[1:0][MemClkDis[5]]) ENDIF. 1=Turn off DLL circuitry.
3:0	Reserved.

D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0] Fence2 Threshold

BIOS: 2.10.5.3.3.

Table 197: Index Mapping for D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0]

D18F2x98_dct[1:0][31:0]	Function
0D0F_2031h	Clock 0
0D0F_2131h	Clock 1
0D0F_2231h	Clock 2
0D0F_8031h	Cmd/Addr 0
0D0F_8131h	Cmd/Addr 1

Table 197: Index Mapping for D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0]

0D0F_8231h	Reserved
0D0F_C031h	Address
0D0F_C131h	Reserved
0D0F_C231h	Reserved

Bits	Description										
31:5	Reserved.										
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="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</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>Description</u>	0h	No delay	1h	1/64 MEMCLK delay	Eh-2h	<Fence2ThresholdTxPad> MEMCLK delay	Fh	15/64 MEMCLK delay
<u>Bits</u>	<u>Description</u>										
0h	No delay										
1h	1/64 MEMCLK delay										
Eh-2h	<Fence2ThresholdTxPad> MEMCLK delay										
Fh	15/64 MEMCLK delay										

D18F2x9C_x0D0F_4009_dct[1:0] Phy Cmp Configuration

Cold reset: 0000_2000h.

Bits	Description										
31:16	Reserved.										
15:14	<p>CmpVioLvl: cmp voltage level. Read-write. BIOS: See 2.10.5.3.1. This field specifies the VDDIO voltage level. Setting this bit in DCT0 adjusts the level for DCT0 and DCT1. Setting this bit in DCT1 has no effect.</p> <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										
13:4	Reserved.										
3:2	<p>ComparatorAdjust: comparator adjust. Read-write. BIOS: D18F2x9C_x0D0F_0[F,8:0]1F_dct[1:0][RxVioLvl]. This field specifies the adjustment signals for the comparator differential amplifier. Setting this bit in DCT0 adjusts the comparator for DCT0 and DCT1. Setting this bit in DCT1 has no effect.</p> <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										
1:0	Reserved.										

D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0] Transmit PreDriver Calibration

Cold reset: xxxx_xxxxh. BIOS: See [2.10.5.3.4](#).

Table 198: [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	<p>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.</p>
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: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, TxPreN]] . 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: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, 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 2.10.5.3.4.

Table 199: 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: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, 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.
5:0	TxPreN: NMOS predriver calibration code. Read-write. This field specifies the falling edge slew rate of the transmit pad. See: Table 177 [Valid Values for D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0][TxPreP, 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_8021_dct[1:0] DLL CS 6 & 7 Timing Control

Cold reset: 0000_0000h. BIOS: 2.10.5.3.3. D18F2x9C_x0000_0004_dct[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][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"> <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 [2.10.5.10](#).

Bits	Description
31:8	Reserved.
7	Add16Tri: MEMADD[16] tri-state. Read-write. Specifies tri-state control for the memory address[16] signal. 1=Signal is tri-stated. 0=Signal is not tri-stated.
6	Reserved.
5	Add17Tri: MEMADD[17] tri-state. Read-write. Specifies tri-state control for the memory address[17] signal. 1=Signal is tri-stated. 0=Signal is not tri-stated.
4:1	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_E003_dct[1:0] Phy Calibration Configuration

Cold reset: 0000_0210h.

Bits	Description
31:15	Reserved.
14	DisAutoComp: disable automatic compensation. Read-write. BIOS: See 2.10.5 and 2.10.5.3.4 . 1=Disable the compensation control state machine. 0=The phy automatic compensation engine is enabled. Setting this bit in DCT0 or DCT1 disables the compensation engine for DCT0 and DCT1.
13	DisablePredriverCal: disable predriver calibration. Read-write. BIOS: See 2.10.5.3.4 . 1=Disables hardware update of predriver calibration codes. Setting this bit in DCT0 affects both DCT0 and DCT1; Setting this bit in DCT1 has no effect.
12:0	Reserved.

D18F2x9C_x0D0F_E006_dct[1:0] Phy PLL Lock Time

Cold reset: 0000_0190h.

Bits	Description
31:16	Reserved.
15:0	PLLLockTime: pll lock time. Read-write. BIOS: 2.10.5.3.2 . Specifies the number of 5ns periods the phy waits for PLLs to lock during a frequency change.

D18F2x9C_x0D0F_E008_dct[1:0] Phy Fence Register

Cold reset: 0000_0013h.

Bits	Description
31:5	Reserved.
4:0	FenceValue: fence value. Read-only. Specifies the fence value used to create the fence bit in the DLL delay registers. See 2.10.5.3.3 [Phy Fence Programming] .

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 2.10.5.3.2 [DRAM Channel Frequency Change] .
3:0	Reserved.

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: 2.10.5.3 . 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.

D18F2xA0_dct[1:0] DRAM Controller MiscellaneousReset: 0000_0000h. Read-only; Updated-by-hardware. See [2.10.1 \[DCT Configuration Registers\]](#).

Bits	Description
31	RcvParErr: parity error status. 1=The M[D:A]_ERR_L pin is asserted. 0=The M[D:A]_ERR_L pin is not asserted. Reserved if D18F2xA8_dct[1:0][LrDimmErrOutMonEn] !=1.
30:0	Reserved.

D18F2xA4 DRAM Controller Temperature Throttle

See [2.10.1 \[DCT Configuration Registers\]](#) and [2.10.10 \[DRAM On DIMM Thermal Management and Power Capping\]](#).

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. MSRC001_0079[BwCapCmdThrottleMode] is an alias of D18F2xA4[BwCapCmdThrottleMode].</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 2.10.5.8 [DRAM Training]) are complete. See MSRC001_0079.</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. 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>Throttle commands by 100%. DRAM devices are placed into power down mode. This mode should only be used to prevent physical damage as system timeouts may occur. Reserved if (D18F2x94_dct[1:0][PowerDownEn]==0 (BwCapEn==1 && BwCapCmdThrottleMode == 1110b)).</td> </tr> </tbody> </table> <p>Throttling should not be enabled until after DRAM initialization (D18F2x110[DramEnable]=1) and training (see 2.10.5.8 [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	Throttle commands by 100%. DRAM devices are placed into power down mode. This mode should only be used to prevent physical damage as system timeouts may occur. Reserved if (D18F2x94_dct[1:0][PowerDownEn]==0 (BwCapEn==1 && BwCapCmdThrottleMode == 1110b)).
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	Throttle commands by 100%. DRAM devices are placed into power down mode. This mode should only be used to prevent physical damage as system timeouts may occur. Reserved if (D18F2x94_dct[1:0][PowerDownEn]==0 (BwCapEn==1 && BwCapCmdThrottleMode == 1110b)).																		
11	<p>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. See MSRC001_0079.</p>																		
10:9	Reserved.																		
8	<p>ODTSEn: on DIMM temperature sensor enable. Read-write. Reset: 0. BIOS: See 2.10.5.6. 1=Enables the monitoring of the EVENT_L pin and indicates whether the on DIMM temperature sensors of the DIMMs on a channel are 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].</p>																		
7:0	Reserved.																		

D18F2xA8_dct[1:0] DRAM Controller Miscellaneous 2

See 2.10.1 [DCT Configuration Registers].

Table 200: BIOS Recommendations for D18F2xA8_dct[1:0][CsMux45]

Condition			D18F2xA8_dct[1:0]
NumDimmSlots	D18F2x80_dct[1:0][Dimm0AddrMap]	D18F2x60_dct[1:0][RankDef]	CsMux45
1,2	7h, 9h	3h	1
	Ah, Bh	2h, 3h	1
1,2,3	-	-	0

Table 201: BIOS Recommendations for D18F2xA8_dct[1:0][CsMux67]

Condition							D18F2xA8_dct[1:0]
NumDimmSlots	D18F2x80_dct[1:0][Dimm0AddrMap]	D18F2x60_dct[1:0][RankDef]	D18F2x80_dct[1:0][Dimm1AddrMap]	D18F2x64_dct[1:0][RankDef]	D18F2x80_dct[1:0][Dimm2AddrMap]	D18F2x68_dct[1:0][RankDef]	CsMux67
2	-	-	7h, 9h	3h	-	-	1
	-	-	Ah, Bh	2h, 3h	-	-	1
3	7h, 9h	3h	-	-	-	-	1
	Ah, Bh	2h, 3h	-	-	-	-	1
	-	-	7h, 9h	3h	-	-	1
	-	-	Ah, Bh	2h, 3h	-	-	1
	-	-	-	-	7h, 9h	3h	1
	-	-	-	-	Ah, Bh	2h, 3h	1
1,2,3	-	-	-	-	-	-	0

Bits	Description
31:30	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. Same-for-all. 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].
27	CsMux67: chip select mux 6 and 7. See: CsMux45. BIOS: Table 201. See also CSTimingMux67.
26	CsMux45: chip select mux 4 and 5. Read-Write. Reset: 0. BIOS: Table 200. 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	<p>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, changing the effective reference point to be earlier than noted in D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0] and D18F2x9C_x0000_00[4A:30]_dct[1:0].</p> <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>1.0 MEMCLK early</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	0 MEMCLK early	01b	0.5 MEMCLK early	10b	1.0 MEMCLK early	11b	Reserved										
Bits	Description																				
00b	0 MEMCLK early																				
01b	0.5 MEMCLK early																				
10b	1.0 MEMCLK early																				
11b	Reserved																				
23	Reserved.																				
22	<p>PrtlChPDEnhEn: partial channel power down enh enable. Read-write. Reset: 0. BIOS: 1. 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.</p>																				
21	<p>AggrPDEn: aggressive power down enable. Read-write. Reset: 0. BIOS: IF (AM3r2) THEN 0 ELSE 1 ENDIF. 1=The DCT places the DRAM devices in precharge power down mode when pages are closed as specified by D18F2x248_dct[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].</p>																				
20:16	Reserved.																				
15:8	<p>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].</p> <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	<p>LrDimmMrsCtrl: LRDIMM MRS control. Read-write. Reset: 0. BIOS: IF (Any D18F2x[6C:60]_dct[1:0][RankDef] >= 10b) THEN 1 ELSE 0 ENDIF. Specifies the buffer MRS control state for hardware initiated MRS commands. 0=Buffer broadcasts MRS commands to all physical ranks. 1=The DCT sets MRS bit 13 to broadcast MRS commands to all physical ranks. See also F0RC14[MRSCoMmandControl].</p>																				
6	<p>LrDimmErrOutMonEn: LRDIMM error out monitor enable. Read-write. Reset: 0. BIOS: See 2.10.5.7.1. 1=Enables the parity error input, ERR, and masks parity errors, preventing it from propagating to the NB MCA logic. Masking begins 20 MEMCLKs after an LRDIMM training start RCW has been sent. 0=Disables the parity error input and parity error masking. See also D18F2x90_dct[1:0][ParEn]. BIOS monitors the error status in D18F2xA0_dct[1:0][RcvParErr].</p>																				
5	<p>SubMemclkRegDly: Sub-one MEMCLK register delay. Read-write. Reset: 0. BIOS: ~D18F2x90_dct[1:0][UnbuffDimm]. 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 Tcwl or Tc1. 0=The delay through the DIMM register and routing delay is at least 1 MEMCLK. The DRAM controller adds 1 MEMCLK to calculate Tcwl and Tc1. Reserved if D18F2x90_dct[1:0][UnbuffDimm].</p>																				

4	ExtendedParityEn: extended parity enable. Read-write. Reset: 0. BIOS: IF (CsMux67 == 1 CsMux45 == 1) THEN 1 ELSE 0 ENDIF. 0=RDIMM parity calculation. 1=Extended parity checking for LRDIMM extended address bits A[17:16]. See also FORC11 [ParityCalculation]. Reserved if ~ D18F2x90_dct[1:0] [ParEn].
3	LrDimmEnhRefEn: LRDIMM enhanced refresh enable. Read-write. Reset: 0. BIOS: LRDIMM . Selects the DCT refresh behavior when D18F2x[6C:60]_dct[1:0] [RankDef] = 1xb. 0=Refresh commands are issued by the DCT to a logical rank and then broadcast by the buffer to all associated physical ranks . 1=Refresh commands are issued by the DCT to a specific physical rank . See also D18F2x228_dct[1:0] [Tstag0]. See also FORC14 [RefreshPrechargeCommandControl].
2	CSTimingMux67: CS timing mux 6 and 7. Read-write. Reset: 0. BIOS: 2.10.5.3.3 . Selects the register used for the setup time of the CS[7:6] pins versus MEMCLK. 0= D18F2x9C_x0000_0004_dct[1:0] [CsOdtSetup]. 1= D18F2x9C_x0000_0004_dct[1:0] [AddrCmd-Setup].
1:0	Reserved.

D18F2xAC DRAM Controller Temperature Status

Cold reset: 0000_0000h.

Bits	Description
31:4	Reserved.
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	<p>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.</p> <table border="1"> <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										

D18F2xFC P-state Power Information 2

Read-only.

Bits	Description
31:24	Reserved.
23:22	PwrDiv7: P7 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
21:20	PwrDiv6: P6 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
19:18	PwrDiv5: P5 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
17:16	PwrDiv4: P4 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
15:14	PwrDiv3: P3 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
13:12	PwrDiv2: P2 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
11:10	PwrDiv1: P1 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
9:8	PwrDiv0: P0 power divisor. See D18F2xF8[PwrValue0] . Value: Product-specific.
7:0	PwrValue4: P4 power value. See D18F2xF8[PwrValue0] . Value: Product-specific.

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\[17C:140,7C:40\]\[DramBase\]](#) must be zero.
- IntLvRgnSwapEn can only be asserted if the probe filter is disabled. (~PrbFltrEn)

- See [D18F2x110\[DctSelIntLvEn\]](#). See [2.10.7 \[Memory Hoisting\]](#).

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; Updated-by-hardware . 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; Updated-by-hardware . 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 2.10.5.7 [DRAM Device and Controller Initialization]) or have exited from self refresh (D18F2x90_dct[1:0][ExitSelfRef] transitions from 1 to 0).																		
7:6	DctSelIntLvAddr: DRAM controller select channel interleave address bit . IF (D18F2x118[Lock-DramCfg]) THEN Read-only. ELSE Read-write. ENDIF. BIOS: 11b. Specifies how interleaving is selected between the DCTs when DctSelIntLvEn =1. In all cases, if the select function is low, then DCT0 is selected; if the select function is high, then DCT1 is selected. The select function is dependent on DctSelIntLvAddr and D18F1x124[DramIntlvEn] as follows: <table border="1" data-bbox="267 1554 1299 1764"> <thead> <tr> <th><u>DctSelIntLvAddr</u></th> <th><u>DramIntlvEn</u></th> <th><u>DCT Select Function</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>xxx</td> <td>Address bit 6.</td> </tr> <tr> <td>01b</td> <td>000</td> <td>Address bit 12.</td> </tr> <tr> <td>01b</td> <td>xx1</td> <td>Reserved.</td> </tr> <tr> <td>10b</td> <td>xxx</td> <td>Hash: exclusive OR of address bits[20:16, 6].</td> </tr> <tr> <td>11b</td> <td>xxx</td> <td>Hash: exclusive OR of address bits[20:16, 9].</td> </tr> </tbody> </table>	<u>DctSelIntLvAddr</u>	<u>DramIntlvEn</u>	<u>DCT Select Function</u>	00b	xxx	Address bit 6.	01b	000	Address bit 12.	01b	xx1	Reserved.	10b	xxx	Hash: exclusive OR of address bits[20:16, 6].	11b	xxx	Hash: exclusive OR of address bits[20:16, 9].
<u>DctSelIntLvAddr</u>	<u>DramIntlvEn</u>	<u>DCT Select Function</u>																	
00b	xxx	Address bit 6.																	
01b	000	Address bit 12.																	
01b	xx1	Reserved.																	
10b	xxx	Hash: exclusive OR of address bits[20:16, 6].																	
11b	xxx	Hash: exclusive OR of address bits[20:16, 9].																	

5	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 when written to DRAM (even bits of the ECC calculation on even beats, odd bits of the ECC calculation on odd beats) such that a dead bit of a DRAM device is correctable. See 2.13.2 [DRAM Considerations for ECC].
4	Reserved.
3	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.
2	DctSelIntLvEn: DRAM controller interleave enable. IF (D18F2x118[LockDramCfg] (D18F5x84[DctEn]!=11b)) THEN Read-only. ELSE Read-write. ENDIF. BIOS: See 2.10.5.6. 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[17C:140,7C: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.
1	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.
0	DctSelHiRngEn: DRAM controller select high range enable. IF (D18F2x118[LockDramCfg] (D18F5x84[DctEn]!=11b)) 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.

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 2.10.7.2 [DctSelBaseOffset Programming].
9:0	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	<p>MctVarPriCntLmt: variable priority time limit. Read-write. Reset: 0000b. BIOS: IF (~MultiLink && (UmaDr UmaIscm) && (D18F2x118[MctPriIsoc]==”Variable”)) THEN “160 ns” ELSE “80 ns” ENDIF.</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>80 ns</td> <td>1000b</td> <td>720 ns</td> </tr> <tr> <td>0001b</td> <td>160 ns</td> <td>1001b</td> <td>800 ns</td> </tr> <tr> <td>0010b</td> <td>240 ns</td> <td>1010b</td> <td>880 ns</td> </tr> <tr> <td>0011b</td> <td>320 ns</td> <td>1011b</td> <td>960 ns</td> </tr> <tr> <td>0100b</td> <td>400 ns</td> <td>1100b</td> <td>1040 ns</td> </tr> <tr> <td>0101b</td> <td>480 ns</td> <td>1101b</td> <td>1120 ns</td> </tr> <tr> <td>0110b</td> <td>560 ns</td> <td>1110b</td> <td>1200 ns</td> </tr> <tr> <td>0111b</td> <td>640 ns</td> <td>1111b</td> <td>1280 ns</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	0000b	80 ns	1000b	720 ns	0001b	160 ns	1001b	800 ns	0010b	240 ns	1010b	880 ns	0011b	320 ns	1011b	960 ns	0100b	400 ns	1100b	1040 ns	0101b	480 ns	1101b	1120 ns	0110b	560 ns	1110b	1200 ns	0111b	640 ns	1111b	1280 ns
Bits	Description	Bits	Description																																		
0000b	80 ns	1000b	720 ns																																		
0001b	160 ns	1001b	800 ns																																		
0010b	240 ns	1010b	880 ns																																		
0011b	320 ns	1011b	960 ns																																		
0100b	400 ns	1100b	1040 ns																																		
0101b	480 ns	1101b	1120 ns																																		
0110b	560 ns	1110b	1200 ns																																		
0111b	640 ns	1111b	1280 ns																																		
27	Reserved.																																				
26:24	<p>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.</p>																																				
23	Reserved.																																				
22:20	<p>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.</p>																																				
19	<p>LockDramCfg. Write-1-only. Reset: 0. BIOS: See 2.10.8 [DRAM CC6 Storage], 2.5.3.3.3 [Core C6 (CC6) State].</p> <p>The following registers are read-only if <code>LockDramCfg=1</code>; 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] • D18F2x250_dct[1:0] [DRAM Loopback and Training Control] • D18F4x128[CoreStateSaveDestNode] • D18F1x[17C:140,7C:40] [DRAM Base/Limit] • D18F1x120 [DRAM Base System Address] • D18F1x124 [DRAM Limit System Address] • D18F2x118[CC6SaveEn] 																																				

18	CC6SaveEn. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0. 1=CC6 save area is enabled. See 2.5.3.7 [BIOS Requirements for Initialization]. BIOS: (D18F4x118[PwrGateEnCstAct0] D18F4x118[PwrGateEnCstAct1] D18F4x11C[PwrGateEnCstAct2]).										
17:16	MctPriScrub: scrubber priority. Read-write. Reset: 00b. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</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>	Bits	Description	00b	Medium	01b	Reserved	10b	High	11b	Variable
Bits	Description										
00b	Medium										
01b	Reserved										
10b	High										
11b	Variable										
15:14	MctPriTrace: trace-mode request priority. See: MctPriCpuRd. Read-write. Reset: 10b. This must be set to high.										
13:12	MctPriIloc: display refresh read priority. See: MctPriCpuRd. Read-write. Reset: 10b. See 2.9.3.2.5 [Display Refresh And IFCM].										
11:10	MctPriWr: default write priority. See: MctPriCpuRd. Read-write. Reset: 01b.										
9:8	MctPriDefault: default non-write priority. See: MctPriCpuRd. Read-write. Reset: 00b.										
7:6	MctPriIlocWr: 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.										
5:4	MctPriIlocRd: 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.										
3:2	MctPriCpuWr: CPU write priority. See: MctPriCpuRd. Read-write. Reset: 01b.										
1:0	MctPriCpuRd: CPU read priority. Read-write. Reset: 00b. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</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>	Bits	Description	00b	Medium	01b	Low	10b	High	11b	Variable
Bits	Description										
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	MctScrubEn: MCT scrub enable. Read-write. Reset: 0. 1=Enables periodic flushing of prefetches and writes based on the DRAM scrub rate. This is used to ensure that prefetch and write data aging is not so long that soft errors accumulate and become uncorrectable. When enabled, each DRAM scrub event causes a single prefetch to be de-allocated (the oldest one) and all queued writes to be flushed to DRAM.
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 2.10.5.8.6 [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	<p>PrefTwoConf: prefetch two-ahead confidence. Read-write. Reset: 011b. BIOS: 011b. Confidence level required in order to prefetch two 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><PrefTwoConf*2></td> </tr> <tr> <td>111b</td> <td>14</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	000b	0	110b-001b	<PrefTwoConf*2>	111b	14		
<u>Bits</u>	<u>Description</u>										
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	<p>PrefConfSat: prefetch confidence saturation. Read-write. Reset: 00b. BIOS: 00b. Specifies the point at which prefetch confidence level saturates and stops incrementing.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	00b	15	01b	7	10b	3	11b	Reserved
<u>Bits</u>	<u>Description</u>										
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: 10h. Specifies the number of writes in the memory controller queue before they are burst into the DCTs.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>32</td> </tr> <tr> <td>1Dh-01h</td> <td><32-MctWrLimit></td> </tr> <tr> <td>1Eh</td> <td>2</td> </tr> <tr> <td>1Fh</td> <td>Write bursting disabled</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00h	32	1Dh-01h	<32-MctWrLimit>	1Eh	2	1Fh	Write bursting disabled
<u>Bits</u>	<u>Description</u>										
00h	32										
1Dh-01h	<32-MctWrLimit>										
1Eh	2										
1Fh	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><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	00b	0	01b	2	10b	4	11b	Reserved
<u>Bits</u>	<u>Description</u>										
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 202: BIOS Recommendations for D18F2x1B0, D18F2x1B4

Condition:	D18F2x1B0	D18F2x1B4	
DdrRate	DcqBwThrotWm	DcqBwThrotWm1	DcqBwThrotWm2
667	0h	3h	4h
800		3h	5h
1066		4h	6h
1333		5h	8h
1600		6h	9h
1866		7h	Ah

Bits	Description																				
31:28	DcqBwThrotWm: dcq bandwidth throttle watermark. Read-write. Reset: 3h. BIOS: Table 202 . Specifies the number of outstanding DRAM read requests before new DRAM prefetch requests and speculative prefetch requests are throttled. 0h=Throttling is determined by D18F2x1B0[DcqBwThrotWm1, DcqBwThrotWm2] . Legal values are 0h through Ch.																				
27:25	PrefFiveConf: prefetch five-ahead confidence. Read-write. Reset: 110b. BIOS: 111b. Confidence level required in order to prefetch five 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><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	PrefFourConf: prefetch four-ahead confidence. Read-write. Reset: 101b. BIOS: 111b. Confidence level required in order to prefetch four 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><PrefFourConf*2></td> </tr> <tr> <td>111b</td> <td>14</td> </tr> </tbody> </table>	Bits	Description	000b	0	110b-001b	<PrefFourConf*2>	111b	14												
Bits	Description																				
000b	0																				
110b-001b	<PrefFourConf*2>																				
111b	14																				
21	Reserved.																				
20	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.																				
19:18	Reserved.																				
17:15	SpecPrefThreshold: speculative prefetch threshold. Read-write. Reset: 111b. Specifies the threshold for the per core hit counters. The counters increment on every hit and decrement on every miss. The counters saturate at 64. The L3 predictor predicts if the per core counter is greater than: <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>Always predict hit.</td> <td>100b</td> <td>32</td> </tr> <tr> <td>001b</td> <td>8</td> <td>101b</td> <td>40</td> </tr> <tr> <td>010b</td> <td>16</td> <td>110b</td> <td>48</td> </tr> <tr> <td>011b</td> <td>24</td> <td>111b</td> <td>56</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	Always predict hit.	100b	32	001b	8	101b	40	010b	16	110b	48	011b	24	111b	56
Bits	Description	Bits	Description																		
000b	Always predict hit.	100b	32																		
001b	8	101b	40																		
010b	16	110b	48																		
011b	24	111b	56																		

14	SpecPrefMis: speculative prefetch predict miss. Read-write. Reset: 0. 1=The L3 predictor always predicts miss for L3 demand reads. This bit overrides the setting in SpecPrefThreshold.												
13	SpecPrefDis: speculative prefetch disable. Read-write. Reset: 0. 1=Disable the L3 speculative miss prefetcher.												
12	EnSplitDctLimits: split DCT write limits enable. Read-write. Reset: 0. BIOS: 1. 1=The number of writes specified by D18F2x11C [DctWrLimit, MctWrLimit] is per DCT. 0=The number of writes specified by D18F2x11C [DctWrLimit, MctWrLimit] is total writes independent of DCT.												
11	DisIoCohPref: disable coherent prefetched for IO. Read-write. Reset: 0. 1=Probes are not generated for prefetches generated for reads from IO devices.												
10:8	CohPrefPrbLmt: coherent prefetch probe limit. Read-write. Reset: 000b. BIOS: 0. Specifies the maximum number of probes that can be outstanding for memory prefetch requests. See 2.9.4.1 [Probe Filter] for additional requirements. <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>16 outstanding probes</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	16 outstanding probes	1xxb	Reserved
<u>Bits</u>	<u>Description</u>												
000b	Probing disabled for memory prefetch requests												
001b	4 outstanding probes												
010b	8 outstanding probes												
011b	16 outstanding probes												
1xxb	Reserved												
7:6	Reserved.												
5:4	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. <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	Reserved.

30:28	S3SmafId: S3 SMAF id. Read-write. Reset: 100b. SMAF encoding of D18F3x[84:80] corresponding to the ACPI S3 state when <code>FlushWrOnS3StpGnt=1</code> . Reserved when <code>FlushWrOnS3StpGnt=0</code> .
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 <code>S3SmafId</code> , indicating entry into the ACPI S3 state. See D18F2xA8_dct[1:0][FastSelfRefEntryDis] , D18F2x11C[FlushWrOnStpGnt] .
26	Reserved.
25:23	Reserved.
22	SpecPrefDisWm1: speculative prefetch disable watermark 1. Read-write. Reset: 0. 0=Disable speculative prefetches at the <code>DcqBwThrotWm2</code> limit. 1=Disable speculative prefetches at the <code>DcqBwThrotWm1</code> 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 <code>DcqBwThrotWm1</code> limit to the <code>DcqBwThrotWm2</code> limit. 0=Requests do not allocate a new entry. 1=Requests allocate a new entry; defined only if (<code>DemandAlloWm1 & DemandPropWm2</code>).
16	DemandPropWm2: demand request propagate watermark 2. Read-write. Reset: 1. Specifies the behavior from the <code>DcqBwThrotWm1</code> limit to the <code>DcqBwThrotWm2</code> limit. 0=Requests do not update existing entries. 1=Requests update existing entries; defined only if (<code>DemandPropWm1=1</code>).
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 <code>DcqBwThrotWm1</code> limit. 0=Requests do not allocate a new entry. 1=Requests allocate a new entry; defined only if (<code>DemandPropWm1=1</code>).
10	DemandPropWm1: demand request propagate watermark 1. Read-write. Reset: 1. Specifies the behavior prior to the <code>DcqBwThrotWm1</code> 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 202 . Specifies a prefetch throttling watermark based on the number of outstanding DRAM read requests. This field is reserved when <code>D18F2x1B0[DcqBwThrotWm] != 0</code> . 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 202 . Specifies a prefetch throttling watermark based on the number of outstanding DRAM read requests. This field is reserved when <code>D18F2x1B0[DcqBwThrotWm] != 0</code> . 0h=Throttling is disabled. Legal values are 0h through 18h.

D18F2x200_dct[1:0] DRAM Timing 0

BIOS: 2.10.5.4. See 2.10.1 [DCT Configuration Registers].

Bits	Description								
31:30	Reserved.								
29:24	Tras: row active strobe. Read-write. Reset: Fh. 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="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>07h-00h</td> <td>Reserved</td> </tr> <tr> <td>28h-08h</td> <td><Tras> clocks</td> </tr> <tr> <td>3Fh-29h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	07h-00h	Reserved	28h-08h	<Tras> clocks	3Fh-29h	Reserved
Bits	Description								
07h-00h	Reserved								
28h-08h	<Tras> clocks								
3Fh-29h	Reserved								
23:21	Reserved.								
20:16	Trp: row precharge time. Read-write. Reset: 5h. 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. <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>13h-02h</td> <td><Trp> clocks</td> </tr> <tr> <td>1Fh-14h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	01h-00h	Reserved	13h-02h	<Trp> clocks	1Fh-14h	Reserved
Bits	Description								
01h-00h	Reserved								
13h-02h	<Trp> clocks								
1Fh-14h	Reserved								
15:13	Reserved.								
12:8	Tred: RAS to CAS delay. Read-write. Reset: 5h. Specifies the time in memory clock cycles from an activate command to a read/write command, both to the same bank. <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>13h-02h</td> <td><Tred> clocks</td> </tr> <tr> <td>1Fh-14h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	01h-00h	Reserved	13h-02h	<Tred> clocks	1Fh-14h	Reserved
Bits	Description								
01h-00h	Reserved								
13h-02h	<Tred> clocks								
1Fh-14h	Reserved								

7:5	Reserved.								
4:0	<p>Tcl: CAS latency. Read-write. Reset: 4h. 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>0Eh-05h</td> <td><Tcl> clocks</td> </tr> <tr> <td>1Fh-0Fh</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	04h-00h	Reserved	0Eh-05h	<Tcl> clocks	1Fh-0Fh	Reserved
<u>Bits</u>	<u>Description</u>								
04h-00h	Reserved								
0Eh-05h	<Tcl> clocks								
1Fh-0Fh	Reserved								

D18F2x204_dct[1:0] DRAM Timing 1

BIOS: See 2.10.5.4. See 2.10.1 [DCT Configuration Registers].

Bits	Description										
31:28	Reserved.										
27:24	<p>Trtp: read CAS to precharge time. Read-write. Reset: 4h. 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>Ah-4h</td> <td><Trtp> clocks</td> </tr> <tr> <td>Fh-Bh</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	3h-0h	Reserved	Ah-4h	<Trtp> clocks	Fh-Bh	Reserved		
<u>Bits</u>	<u>Description</u>										
3h-0h	Reserved										
Ah-4h	<Trtp> clocks										
Fh-Bh	Reserved										
23:22	Reserved.										
21:16	<p>FourActWindow: four bank activate window. Read-write. Reset: 0. 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>2Ah-06h</td> <td><FourActWindow> clocks</td> </tr> <tr> <td>3Fh-2Bh</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00h	No tFAW window restriction	05h-01h	Reserved	2Ah-06h	<FourActWindow> clocks	3Fh-2Bh	Reserved
<u>Bits</u>	<u>Description</u>										
00h	No tFAW window restriction										
05h-01h	Reserved										
2Ah-06h	<FourActWindow> clocks										
3Fh-2Bh	Reserved										
15:12	Reserved.										
11:8	<p>Trrd: row to row delay (or RAS to RAS delay). Read-write. Reset: 4h. Specifies the minimum time in memory clock cycles between activate commands to different chip select banks.</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>9h-1h</td> <td><Trrd> clocks</td> </tr> <tr> <td>Fh-Ah</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	Reserved	9h-1h	<Trrd> clocks	Fh-Ah	Reserved		
<u>Bits</u>	<u>Description</u>										
0h	Reserved										
9h-1h	<Trrd> clocks										
Fh-Ah	Reserved										
7:6	Reserved.										
5:0	<p>Trc: row cycle time. Read-write. Reset: Bh. 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.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>09h-00h</td> <td>Reserved</td> </tr> <tr> <td>38h-0Ah</td> <td><Trc> clocks</td> </tr> <tr> <td>3Fh-39h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	09h-00h	Reserved	38h-0Ah	<Trc> clocks	3Fh-39h	Reserved		
<u>Bits</u>	<u>Description</u>										
09h-00h	Reserved										
38h-0Ah	<Trc> clocks										
3Fh-39h	Reserved										

D18F2x208_dct[1:0] DRAM Timing 2

BIOS: See 2.10.5.4. See 2.10.1 [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	<p>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.</p> <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>90 ns (all speeds, 512 Mbit)</td> </tr> <tr> <td>010b</td> <td>110 ns (all speeds, 1 Gbit)</td> </tr> <tr> <td>011b</td> <td>160 ns (all speeds, 2 Gbit)</td> </tr> <tr> <td>100b</td> <td>300 ns (all speeds, 4 Gbit)</td> </tr> <tr> <td>101b</td> <td>350 ns (all speeds, 8 Gbit)</td> </tr> <tr> <td>111b-110b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	Reserved	001b	90 ns (all speeds, 512 Mbit)	010b	110 ns (all speeds, 1 Gbit)	011b	160 ns (all speeds, 2 Gbit)	100b	300 ns (all speeds, 4 Gbit)	101b	350 ns (all speeds, 8 Gbit)	111b-110b	Reserved
Bits	Description																
000b	Reserved																
001b	90 ns (all speeds, 512 Mbit)																
010b	110 ns (all speeds, 1 Gbit)																
011b	160 ns (all speeds, 2 Gbit)																
100b	300 ns (all speeds, 4 Gbit)																
101b	350 ns (all speeds, 8 Gbit)																
111b-110b	Reserved																

D18F2x20C_dct[1:0] DRAM Timing 3

See 2.10.1 [DCT Configuration Registers].

Table 203: BIOS Recommendations for D18F2x20C_dct[1:0][Tcwl]

Condition	D18F2x20C_dct[1:0]
DdrRate	Tcwl
667, 800	5
1066	6
1333	7
1600	8
1866	9

Bits	Description
31:12	Reserved.

11:8	Twtr: internal DRAM write to read command delay. Read-write. Reset: 4h. BIOS: See 2.10.5.4. 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. <u>Bits</u> 3h-0h 9h-4h Fh-Ah	<u>Description</u> Reserved <Twtr> clocks Reserved
7:5	Reserved.	
4:0	Tcwl: CAS write latency. Read-write. Reset: 5h. BIOS: Table 203. Specifies the number of memory clock cycles from internal write command to first write data in at the DRAM. <u>Bits</u> 04h-00h 0Ah-05h 1Fh-0Bh	<u>Description</u> Reserved <Tcwl> clocks Reserved

D18F2x210_dct[1:0]_nbp[3:0] DRAM NB P-state

See 2.10.1 [DCT Configuration Registers]. For D18F2x210_dct[1:0]_nbp[x], x=D18F1x10C[NbPsSel]; see D18F1x10C[NbPsSel].

Table 204: Register Mapping for D18F2x210_dct[1:0]_nbp[3:0]

Register	Function
D18F2x210_dct[1:0]_nbp0	P-state 0
D18F2x210_dct[1:0]_nbp1	P-state 1
D18F2x210_dct[1:0]_nbp2	P-state 2
D18F2x210_dct[1:0]_nbp3	P-state 3

Table 205: BIOS Recommendations for D18F2x210_dct[1:0]_nbp[3:0][DataTxFifoWrDly]

Condition	D18F2x210_dct[1:0]_nbp[3:0]
D18F2x210_dct[1:0]_nbp[3:0][RdPtrInit]	NbPsSel
0100b	D18F1x10C[NbPsSel]==0 D18F1x10C[NbPsSel]>0
0101b	-
0110b	-
	DataTxFifoWrDly
	010b
	001b
	000b

Table 206: BIOS Recommendations for D18F2x210_dct[1:0]_nbp[3:0][RdPtrInit]

Condition	D18F2x210_dct[1:0]_nbp[3:0]
DdrRate	RdPtrInit
667, 800, 1066, 1333, 1600	0110b
1866	0100b

Bits	Description
------	-------------

31:22	<p>MaxRdLatency: maximum read latency. Read-write. Reset: 000h. BIOS: 2.10.5.8.5. 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><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	000h	0 NCLKs	3FEh-001h	<MaxRdLatency> NCLKs	3FFh	1023 NCLKs										
<u>Bits</u>	<u>Description</u>																		
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: Table 205. Specifies the DCT to phy write data FIFO delay.</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 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>	<u>Bits</u>	<u>Description</u>	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
<u>Bits</u>	<u>Description</u>																		
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 206. 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. See also 2.10.5.2 [NB P-state Specific Configuration].</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0011b-0000b</td> <td>Reserved</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>	<u>Bits</u>	<u>Description</u>	0011b-0000b	Reserved	0100b	2 MEMCLKs	0101b	1.5 MEMCLKs	0110b	1 MEMCLK	1111b-0111b	Reserved						
<u>Bits</u>	<u>Description</u>																		
0011b-0000b	Reserved																		
0100b	2 MEMCLKs																		
0101b	1.5 MEMCLKs																		
0110b	1 MEMCLK																		
1111b-0111b	Reserved																		

D18F2x214_dct[1:0] DRAM Timing 4

Bits	Description
31:20	Reserved.

19:16	<p>TwrrwrSdSc: write to write timing same DIMM same chip select. Read-write. Reset: 1h. BIOS: See 2.10.5.5.2 [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. For LRDIMMs this term applies to accesses to the same physical rank.</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. For LRDIMMs this term applies to accesses to a different physical rank.												
7:4	Reserved.												
3:0	<p>TwrrwrDd: write to write timing different DIMM. Read-write. Reset: 2h. BIOS: See 2.10.5.5.2 [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. For LRDIMMs this term applies to accesses to a different DIMM.</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] DRAM Timing 5

See 2.10.1 [DCT Configuration Registers].

Bits	Description								
31:28	Reserved.								
27:24	<p>TrdrdSdSc: read to read timing same DIMM same chip select. Read-write. Reset: 1h. BIOS: See 2.10.5.5.1 [TrdrdSdSc, TrdrdSdDc, 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. For LRDIMMs this term applies to accesses to the same physical rank.</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. For LRDIMMs this term applies to accesses to a different physical rank.								
15:12	Reserved.								

11:8	<p>Twrrd: write to read DIMM termination turnaround. Read-write. Reset: 2h. BIOS: See 2.10.5.5.3 [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. For LRDIMMs this term applies to accesses to a different physical rank or a different DIMM.</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. Reset: 3h. BIOS: See 2.10.5.5.1 [TrdrdSdSc, TrdrdSdDc, 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. For LRDIMMs this term applies to accesses to a different DIMM.</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] DRAM Timing 6

See [2.10.1 \[DCT Configuration Registers\]](#).

Bits	Description								
31:21	Reserved.								
20:16	<p>TrwtWB: read to write turnaround for opportunistic write bursting. Read-write. Reset: 4h. 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. Reset: 3h. BIOS: See 2.10.5.5.4 [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

See 2.10.1 [DCT Configuration Registers].

Table 207: BIOS Recommendations for D18F2x220_dct[1:0][Tmod]

Condition	D18F2x220_dct[1:0]
DdrRate	Tmod
667, 800, 1066, 1333, 1600	Ch
1866	Eh

Bits	Description								
31:13	Reserved.								
12:8	<p>Tmod: mode register command delay. Read-write. Reset: Ch. BIOS: Table 207. Specifies the minimum time in memory clock cycles from an MRS command to another non-MRS command (excluding NOP and DES).</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>12h-2h</td> <td><Tmod> clocks</td> </tr> <tr> <td>1Fh-13h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	1h-0h	Reserved	12h-2h	<Tmod> clocks	1Fh-13h	Reserved
Bits	Description								
1h-0h	Reserved								
12h-2h	<Tmod> clocks								
1Fh-13h	Reserved								
7:4	Reserved.								
3:0	<p>Tmrd: mode register command cycle time. Read-write. Reset: 4h. BIOS: IF (LRDIMM) THEN 6 ELSE 4 ENDIF. 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

See 2.10.1 [DCT Configuration Registers].

Table 208: BIOS Recommendations for D18F2x224_dct[1:0][Tzqcs, Tzqoper]

Condition	D18F2x224_dct[1:0]	
	Tzqcs	Tzqoper
667, 800, 1066, 1333, 1600	100b	1000b
1866	101b	1010b

Bits	Description
31:11	Reserved.

10:8	Tzqcs: Zq short cal command delay. Read-write. Reset: 4h. BIOS: Table 208 . Specifies the minimum time in memory clock cycles from a ZQCS command to any other command (excluding NOP and DES) on the channel.			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	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	Tzqoper: Zq long cal command delay. Read-write. Reset: 8h. BIOS: Table 208 . Specifies the minimum time in memory clock cycles from a ZQCL command to any other command (excluding NOP and DES) on the channel.			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	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 [2.10.1 \[DCT Configuration Registers\]](#).

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	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 . See also D18F2xA8_dct[1:0][LrDimmEnhRef] and F1RC9 , F1RC10
	<u>Bits</u> <u>Description</u>
	00h 0 clocks
	FEh-01h <Tstag0> clocks
	FFh 255 clocks

D18F2x22C_dct[1:0] DRAM Timing 10

Reset: 0000_000Ch. See [2.10.1 \[DCT Configuration Registers\]](#).

Bits	Description
------	-------------

31:5	Reserved.																												
4:0	<p>Twr: write recovery. Read-write. BIOS: See 2.10.5.4. Specifies the minimum time from the last data write until the chip select bank precharge.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>4h-0h</td> <td>Reserved</td> </tr> <tr> <td>5h</td> <td>5 clocks</td> </tr> <tr> <td>7h-6h</td> <td><Twr> clocks</td> </tr> <tr> <td>8h</td> <td>8 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>1Fh-11h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	4h-0h	Reserved	5h	5 clocks	7h-6h	<Twr> clocks	8h	8 clocks	9h	Reserved	Ah	10 clocks	Bh	Reserved	Ch	12 clocks	Dh	Reserved	Eh	14 clocks	Fh	Reserved	10h	16 clocks	1Fh-11h	Reserved
Bits	Description																												
4h-0h	Reserved																												
5h	5 clocks																												
7h-6h	<Twr> clocks																												
8h	8 clocks																												
9h	Reserved																												
Ah	10 clocks																												
Bh	Reserved																												
Ch	12 clocks																												
Dh	Reserved																												
Eh	14 clocks																												
Fh	Reserved																												
10h	16 clocks																												
1Fh-11h	Reserved																												

D18F2x[234:230]_dct[1:0] DRAM Read ODT Pattern

Reset: 0000_0000h. BIOS: 2.10.5.5.5. See 2.10.1 [DCT Configuration Registers]. This register is used by BIOS to specify the state of the ODT pins during DDR reads.

Table 209: Register Mapping for D18F2x[234:230]_dct[1:0]

Register	Function
D18F2x230_dct[1:0]	Chip Selects [3:0]
D18F2x234_dct[1:0]	Chip Selects [7:4]

Table 210: Field Mapping for D18F2x[234:230]_dct[1:0]

Register	Bits			
	31:24	23:16	15:8	7:0
D18F2x230_dct[1:0]	Cs3	Cs2	Cs1	Cs0
D18F2x234_dct[1:0]	Cs7	Cs6	Cs5	Cs4

Bits	Description
31:28	Reserved.
27:24	RdOdtPat: read ODT pattern chip select [7,3]. See: D18F2x[234:230]_dct[1:0][3:0].
23:20	Reserved.
19:16	RdOdtPat: read ODT pattern chip select [6,2]. See: D18F2x[234:230]_dct[1:0][3:0].
15:12	Reserved.
11:8	RdOdtPat: read ODT pattern chip select [5,1]. See: D18F2x[234:230]_dct[1:0][3:0].

7:4	Reserved.
3:0	RdOdtPat: 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

Reset: 0000_0000h. BIOS: 2.10.5.5.5. See 2.10.1 [DCT Configuration Registers]. This register is used by BIOS to specify the state of the ODT pins during DDR writes.

Table 211: Register Mapping for D18F2x[23C:238]_dct[1:0]

Register	Function
D18F2x238_dct[1:0]	Chip Selects [3:0]
D18F2x23C_dct[1:0]	Chip Selects [7:4]

Table 212: Field Mapping for D18F2x[23C:238]_dct[1:0]

Register	Bits			
	31:24	23:16	15:8	7:0
D18F2x238_dct[1:0]	Cs3	Cs2	Cs1	Cs0
D18F2x23C_dct[1:0]	Cs7	Cs6	Cs5	Cs4

Bits	Description
31:28	Reserved.
27:24	WrOdtPat: write ODT pattern chip select [7,3] . See: D18F2x[23C:238]_dct[1:0][3:0].
23:20	Reserved.
19:16	WrOdtPat: write ODT pattern chip select [6,2] . See: D18F2x[23C:238]_dct[1:0][3:0].
15:12	Reserved.
11:8	WrOdtPat: write ODT pattern chip select [5,1] . See: D18F2x[23C:238]_dct[1:0][3:0].
7:4	Reserved.
3:0	WrOdtPat: 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] DRAM ODT Control

Reset: 0000_0000h. See 2.10.1 [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 writes. <table border="0" style="margin-left: 20px;"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0 clocks. (Don't assert ODT)</td> </tr> <tr> <td>1h</td> <td>1 clock</td> </tr> <tr> <td>6h-2h</td> <td><WrOdtOnDuration> clocks</td> </tr> <tr> <td>7h</td> <td>7 clocks</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	0 clocks. (Don't assert ODT)	1h	1 clock	6h-2h	<WrOdtOnDuration> clocks	7h	7 clocks
<u>Bits</u>	<u>Description</u>										
0h	0 clocks. (Don't assert ODT)										
1h	1 clock										
6h-2h	<WrOdtOnDuration> clocks										
7h	7 clocks										

11	Reserved.								
10:8	<p>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.</p> <table border="1"> <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)								
7:4	<p>RdOdtOnDuration: Read ODT On Duration. Read-write. BIOS: IF LRDIMM THEN See 2.10.5.4.1 ELSE 6 ENDIF. Specifies the number of memory clock cycles that ODT is asserted for read.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0 clocks (Don't assert ODT)</td> </tr> <tr> <td>9h-1h</td> <td><RdOdtOnDuration> clocks</td> </tr> <tr> <td>Fh-Ah</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	0h	0 clocks (Don't assert ODT)	9h-1h	<RdOdtOnDuration> clocks	Fh-Ah	Reserved.
Bits	Description								
0h	0 clocks (Don't assert ODT)								
9h-1h	<RdOdtOnDuration> clocks								
Fh-Ah	Reserved.								
3:0	<p>RdOdtTrnOnDly: Read ODT Turn On Delay. Read-write. BIOS: IF LRDIMM THEN See 2.10.5.4.1 ELSE MAX(0, D18F2x200_dct[1:0][Tcl] - D18F2x20C_dct[1:0][Tcwl]) ENDIF. Specifies the number of clock cycles that ODT assertion is delayed relative to read CAS.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0 clocks</td> </tr> <tr> <td>Fh-0h</td> <td><RdOdtTrnOnDly> clocks, Reserved if (RdOdtOnDuration=0)</td> </tr> </tbody> </table>	Bits	Description	0h	0 clocks	Fh-0h	<RdOdtTrnOnDly> clocks, Reserved if (RdOdtOnDuration=0)		
Bits	Description								
0h	0 clocks								
Fh-0h	<RdOdtTrnOnDly> clocks, Reserved if (RdOdtOnDuration=0)								

D18F2x244_dct[1:0] DRAM Controller Miscellaneous 3

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers].

Bits	Description										
31:4	Reserved.										
3:0	<p>PrtlChPDDynDly: partial channel power down dynamic delay. Read-write. BIOS: 2h. Specifies the channel idle hysteresis for fast exit/slow exit mode changes when D18F2xA8_dct[1:0][PrtlChP-DEnhEn]=1.</p> <table border="1"> <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><PrtlChPDDynDly*32> clocks</td> </tr> <tr> <td>8h</td> <td>256 clocks</td> </tr> <tr> <td>Fh-9h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	0 clocks	7h-1h	<PrtlChPDDynDly*32> clocks	8h	256 clocks	Fh-9h	Reserved
Bits	Description										
0h	0 clocks										
7h-1h	<PrtlChPDDynDly*32> clocks										
8h	256 clocks										
Fh-9h	Reserved										

D18F2x248_dct[1:0] DRAM Power Management 0

See 2.10.1 [DCT Configuration Registers].

Table 213: BIOS Recommendations for D18F2x248_dct[1:0][Txp] (SODIMM | UDIMM)

Condition	D18F2x248_dct[1:0]
DdrRate	Txp
667, 800	3h
1066, 1333	4h
1600	5h
1866	6h

Table 214: BIOS Recommendations for D18F2x248_dct[1:0][Txp] (RDIMM | LRDIMM)

Condition		D18F2x248_dct[1:0]
DdrRate	DdrVDDIO	Txp
667, 800	1.5, 1.35	3h
667, 800	1.25	4h
1066, 1333	-	4h
1600	-	5h
1866	-	6h

Table 215: BIOS Recommendations for D18F2x248_dct[1:0][Txpdll]

Condition	D18F2x248_dct[1:0]
DdrRate	Txpdll
667, 800	Ah
1066	Dh
1333	10h
1600	14h
1866	17h

Bits	Description										
31:30	Reserved.										
29:24	<p>AggrPDDelay: aggressive power down delay. Read-Write. Reset: 0. BIOS: 0h. 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].</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><AggrPDDelay> clocks</td> </tr> <tr> <td>3Fh</td> <td>63 clocks</td> </tr> </tbody> </table>	Bits	Description	00h	64 clocks	01h	1 clock	3Eh-02h	<AggrPDDelay> clocks	3Fh	63 clocks
Bits	Description										
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. Reset: 0. BIOS: IF (D18F2xA8_dct[1:0][AggrPDEn]) THEN 01h ELSE 00h ENDIF. Specifies the power down entry delay. If D18F2xA8_dct[1:0][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. Reset: 0Ah. BIOS: Table 215. Specifies the minimum time that the DCT waits to issue a command requiring the DRAM DLL after exiting precharge powerdown mode if the DRAM DLL was disabled.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	09h-00h	Reserved	1Dh-0Ah	<Txpdll> clocks	1Fh-1Eh	Reserved
<u>Bits</u>	<u>Description</u>								
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. Reset: 3h. BIOS: Table 213, Table 214. Specifies the minimum time that the DCT waits to issue a command after exiting precharge powerdown mode.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></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>	<u>Bits</u>	<u>Description</u>	2h-0h	Reserved	8h-3h	<Txp> clocks	Fh-9h	Reserved
<u>Bits</u>	<u>Description</u>								
2h-0h	Reserved								
8h-3h	<Txp> clocks								
Fh-9h	Reserved								

D18F2x24C_dct[1:0] DRAM Power Management 1

See [2.10.1 \[DCT Configuration Registers\]](#).

Table 216: BIOS Recommendations for D18F2x24C_dct[1:0][Tcksrx, Tcksre, Tpd]

Condition	D18F2x24C_dct[1:0]		
	Tcksrx	Tcksre	Tpd
667, 800	5h		3
1066	6h		3
1333	7h		4
1600	8h		4
1866	Ah		5

Bits	Description								
31:30	Reserved.								
29:24	<p>Tcksrx: clock stable to self refresh exit delay. Read-write. Reset: 00h. BIOS: Table 216. 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><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>01h-00h</td> <td>Reserved</td> </tr> <tr> <td>0Eh-02h</td> <td><Tcksrx> clocks</td> </tr> <tr> <td>3Fh-0Fh</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	01h-00h	Reserved	0Eh-02h	<Tcksrx> clocks	3Fh-0Fh	Reserved
<u>Bits</u>	<u>Description</u>								
01h-00h	Reserved								
0Eh-02h	<Tcksrx> clocks								
3Fh-0Fh	Reserved								
23:22	Reserved.								
21:16	<p>Tcksre: self refresh to command delay. Read-write. Reset: 14h. BIOS: Table 216. 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><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>04h-00h</td> <td>Reserved</td> </tr> <tr> <td>27h-05h</td> <td><Tcksre> clocks</td> </tr> <tr> <td>3Fh-28h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	04h-00h	Reserved	27h-05h	<Tcksre> clocks	3Fh-28h	Reserved
<u>Bits</u>	<u>Description</u>								
04h-00h	Reserved								
27h-05h	<Tcksre> clocks								
3Fh-28h	Reserved								

15:14	Reserved.								
13:8	<p>Tckesr: self refresh to command delay. Read-write. Reset: 08h. BIOS: Tpd + 1. Specifies the minimum time in memory clock cycles that the DCT waits to issue a command after entering self refresh.</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>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	<p>Tpd: minimum power down entry to exit. Read-write. Reset: 3h. BIOS: Table 216.</p> <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 [2.10.1 \[DCT Configuration Registers\]](#). See [2.10.5.8.6.1 \[DRAM Training Pattern Generation\]](#).

Bits	Description										
31:14	Reserved.										
13	<p>LfsrRollOver: LFSR roll over. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Specifies the behavior of DataPrbsSeed and the data comparison logic if the generated address wraps around to equal D18F2x25[8,4]_dct[1:0][TgtAddress]. 0=The LFSR will not be re-seeded. 1=The LFSR will be re-seeded.</p>										
12	<p>CmdSendInProg: command in progress. Read-only. 0=DCT is idle. 1=DCT is busy.</p>										
11	<p>SendCmd: send command. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. 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.</p>										
10	<p>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)).</p>										
9:8	<p>CmdTgt: command target. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Specifies the SendCmd command target address mode. See D18F2x25[8,4]_dct[1:0].</p> <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	<p>CmdType: command type. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Specifies the SendCmd command type.</p> <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. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. 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. IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read; write-1-only; cleared-by-hardware. ENDIF. 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 Base

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [DRAM Training Pattern Generation].

Table 217: Register Mapping for D18F2x25[8,4]_dct[1:0]

Register	Function
D18F2x254_dct[1:0]	Target A
D18F2x258_dct[1:0]	Target B

Bits	Description								
31:27	Reserved.								
26:24	TgtChipSelect: target chip select. Read-write. Specifies the chip select. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>CS 0</td> </tr> <tr> <td>110b-001b</td> <td>CS <TgtChipSelect></td> </tr> <tr> <td>111b</td> <td>CS 7</td> </tr> </tbody> </table>	Bits	Description	000b	CS 0	110b-001b	CS <TgtChipSelect>	111b	CS 7
Bits	Description								
000b	CS 0								
110b-001b	CS <TgtChipSelect>								
111b	CS 7								
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 D18F2x250_dct[1:0][CmdType] = 00xb or if D18F2x250_dct[1:0][CmdType] = 010b and D18F2x250_dct[1:0][CmdTgt] = 01b. TgtAddress[9:3] is incremented by one, with wrap around, after each command pair if D18F2x250_dct[1:0][CmdType] = 010b and D18F2x250_dct[1:0][CmdTgt] = 00b.								

D18F2x260_dct[1:0] DRAM Command 1

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [DRAM Training Pattern Generation].

Bits	Description
------	-------------

31:21	Reserved.						
20:0	<p>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>.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved.</td> </tr> <tr> <td>1F_FFFFh-1h</td> <td><CmdCount> commands</td> </tr> </tbody> </table>	Bits	Description	0h	Reserved.	1F_FFFFh-1h	<CmdCount> commands
Bits	Description						
0h	Reserved.						
1F_FFFFh-1h	<CmdCount> commands						

D18F2x264_dct[1:0] DRAM Status 0

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [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>, <code>D18F2x278_dct[1:0]</code>, and <code>D18F2x27C_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_FFDDh-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_FFDDh-1h	<ErrCnt> errors	1FF_FFFEh	1FF_FFFEh errors	1FF_FFFFh	1FF_FFFFh or more errors						
Bits	Description																
0h	0 errors																
1FF_FFDDh-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 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [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[(<code><NibbleErrSts>*4</code>)+3:<code><NibbleErrSts>*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><NibbleErrSts>*4</code>)+3: <code><NibbleErrSts>*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><NibbleErrSts>*4</code>)+3: <code><NibbleErrSts>*4</code>]														
[15]	Data[63:60]														
[16]	ECC[3:0]														
[17]	ECC[7:4]														

D18F2x26C_dct[1:0] DRAM Status 2

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [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 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [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.</p>

D18F2x274_dct[1:0] DRAM DQ Mask Low

See `D18F1x10C[DctCfgSel]`. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [DRAM Training Pattern Generation].

Bits	Description								
31:0	DQMask[31:0]: DQ mask. Read-write. DQMask[63:0] = {D18F2x278_dct[1:0][DQMask[63:32]], DQMask[31:0]}. Reset: 0000_0000_0000_0000h. 1=The corresponding DQ bit will not be compared. 0=The corresponding DQ bit will be compared. See D18F2x264_dct[1:0][ErrCnt].								
	<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[<DQMask>]</td> </tr> <tr> <td>[63]</td> <td>Data[63]</td> </tr> </tbody> </table>	Bit	Description	[0]	Data[0]	[62:1]	Data[<DQMask>]	[63]	Data[63]
Bit	Description								
[0]	Data[0]								
[62:1]	Data[<DQMask>]								
[63]	Data[63]								

D18F2x278_dct[1:0] DRAM DQ Mask High

Bits	Description
31:0	DQMask[63:32]: DQ mask. See: D18F2x274_dct[1:0][DQMask[31:0]].

D18F2x27C_dct[1:0] DRAM ECC Mask

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [DRAM Training Pattern Generation].

Bits	Description								
31:8	Reserved.								
7:0	EccMask: ECC mask. Read-write. 1=The corresponding ECC DQ bit will not be compared. 0=The corresponding ECC DQ bit will be compared. See D18F2x264_dct[1:0][ErrCnt].								
	<table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>ECC[0]</td> </tr> <tr> <td>[6:1]</td> <td>ECC[<EccMask>]</td> </tr> <tr> <td>[7]</td> <td>ECC[7].</td> </tr> </tbody> </table>	Bit	Description	[0]	ECC[0]	[6:1]	ECC[<EccMask>]	[7]	ECC[7].
Bit	Description								
[0]	ECC[0]								
[6:1]	ECC[<EccMask>]								
[7]	ECC[7].								

D18F2x28C_dct[1:0] DRAM Command 2

Reset: 0000_0000h. See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [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 LRDIMMs, BIOS programs this field with the logical rank CS. <table border="0"> <tr> <td><u>Bit</u></td> <td><u>Description</u></td> </tr> <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> </table>	<u>Bit</u>	<u>Description</u>	[0]	CS0	[6:1]	CS<CmdChipSelect>	[7]	CS7
<u>Bit</u>	<u>Description</u>								
[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 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [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="0"> <tr> <td><u>Bits</u></td> <td><u>Description</u></td> </tr> <tr> <td>7h-0h</td> <td><ErrBeatNum> beat</td> </tr> </table>	<u>Bits</u>	<u>Description</u>	7h-0h	<ErrBeatNum> beat		
<u>Bits</u>	<u>Description</u>						
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="0"> <tr> <td><u>Bits</u></td> <td><u>Description</u></td> </tr> <tr> <td>0</td> <td>No error</td> </tr> <tr> <td>1F_FFFFh-1h</td> <td><ErrCmdNum> command</td> </tr> </table>	<u>Bits</u>	<u>Description</u>	0	No error	1F_FFFFh-1h	<ErrCmdNum> command
<u>Bits</u>	<u>Description</u>						
0	No error						
1F_FFFFh-1h	<ErrCmdNum> command						

D18F2x294_dct[1:0] DRAM Status 4

See 2.10.1 [DCT Configuration Registers]. See 2.10.5.8.6.1 [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="0"> <tr> <td><u>Bit</u></td> <td><u>Description</u></td> </tr> <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> </table>	<u>Bit</u>	<u>Description</u>	[0]	Data[0]	[62:1]	Data[<DQErr>]	[63]	Data[63]
<u>Bit</u>	<u>Description</u>								
[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]] .

D18F2x29C_dct[1:0] DRAM Status 6

Reset: 0000_0000h. See [2.10.1 \[DCT Configuration Registers\]](#). See [2.10.5.8.6.1 \[DRAM Training Pattern Generation\]](#).

Bits	Description								
31:8	Reserved.								
7:0	EccErr: ECC error. Read-only. 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][Reset-AllErr] . <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>ECC[0]</td> </tr> <tr> <td>[6:1]</td> <td>ECC[<EccErr>]</td> </tr> <tr> <td>[7]</td> <td>ECC[7].</td> </tr> </tbody> </table>	Bit	Description	[0]	ECC[0]	[6:1]	ECC[<EccErr>]	[7]	ECC[7].
Bit	Description								
[0]	ECC[0]								
[6:1]	ECC[<EccErr>]								
[7]	ECC[7].								

3.6 Device [1F:18]h Function 3 Configuration Registers

See [3.1 \[Register Descriptions and Mnemonics\]](#). See [2.8 \[Configuration Space\]](#).

D18F3x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1603h.
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	See: MSR0000_0410 [31:0].

D18F3x44 MCA NB Configuration

Same-for-all. 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=Enable 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	<p>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, and 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.</p> <p>Note:</p> <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 281 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	<p>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.</p>
25	<p>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.</p>
24	<p>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.</p>
23	<p>ChipKillEccCap: chip-kill ECC mode. Read-only; updated-by-hardware. Reset: 0. 1=Chipkill ECC mode capable; ECC checking is based on x8 ECC symbols (D18F3x180[EccSymbolSize]) and can be used for chipkill. 0=Chipkill ECC mode not capable; ECC checking is based on two interleaved, ungangned 64/8-bit data/ECC lines and x4 ECC symbols and cannot be used for chipkill. See 2.13.2 [DRAM Considerations for ECC].</p>
22	<p>DramEccEn: DRAM ECC enable. Read-write. Reset: 0. 1=Enables ECC check/correct mode. This bit must be set in order for ECC checking/correcting by the NB to be enabled. If set, ECC is checked and correctable errors are corrected irrespective of whether machine check ECC reporting is enabled. The hardware only allows values to be programmed into this field which are consistent with the ECC capabilities of the device as specified in D18F3xE8 [Northbridge Capabilities]. Attempts to write values inconsistent with the capabilities results in this field not being updated. This bit does not affect ECC checking in the northbridge arrays.</p>
21	<p>SyncFloodOnAnyUcErr: sync flood on any UC error. Read-write. Reset: 0. BIOS: 1. 1=Enable sync flood of all links with sync packets on detection of any NB MCA error that is uncorrectable, including northbridge array errors and link protocol errors.</p>
20	<p>SyncFloodOnWDT: sync flood on watchdog timer error. Read-write. Reset: 0. BIOS: 1. 1=Enable sync flood of all links with sync packets on detection of a watchdog timer error.</p>

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	<p>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. See D18F0x150[ForceErrType]. See D18F0x[14C:130][ForceRetryError] for generating errors in retry mode.</p>										
16	<p>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. See D18F0x150[ForceErrType]. See D18F0x[14C:130][ForceRetryError] for generating errors in retry mode.</p>										
15:14	<p>GenLinkSel: link select for CRC error generation. Read-Write. Reset: 00b. Selects the link to be used for CRC error injection through GenCrcErrByte1/GenCrcErrByte0.</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										
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[3:0] = {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. 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. BIOS: 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. BIOS: 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 D18F0x[E4,C4,A4,84] [Link Control].</p>																								
2	<p>SyncFloodOnDramUcEcc: sync flood on uncorrectable DRAM ECC error. Read-write. Reset: 0. BIOS: 1. 1=Enable sync flood of all links with sync packets on detection of an uncorrectable DRAM 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	See: MSR0000_0411 [31:0].

D18F3x4C MCA NB Status High

Bits	Description
31:0	See: MSR0000_0411 [63:32].

D18F3x50 MCA NB Address Low

Bits	Description
31:0	See: MSR0000_0412 [31:0].

D18F3x54 MCA NB Address High

Bits	Description
31:0	See: MSR0000_0412 [63:32].

D18F3x58 Scrub Rate Control

This register specifies the ECC sequential scrubbing rate for lines of memory and cache. See 2.9.5 [Memory Scrubbers]. Accesses to this register must first set [D18F1x10C](#)[DctCfgSel]=0; Accesses to this register with [D18F1x10C](#)[DctCfgSel]=1 are undefined; See erratum 505. Scrub rates are a platform consideration. See 2.13.1.8 [Scrub Rate Considerations].

Bits	Description
31:29	Reserved.

28:24	L3Scrub: L3 cache scrub rate. Read-write. Reset: 00000b. BIOS: 2.10.5.6. Specifies time between 64 B scrub events. See 2.9.4.1 [Probe Filter].			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	00h	Disable sequential scrubbing	10h	1.31 ms
	01h	40 ns ¹	11h	2.62 ms
	02h	80 ns ¹	12h	5.24 ms
	03h	160 ns ¹	13h	10.49 ms
	04h	320 ns ¹	14h	20.97 ms
	05h	640 ns	15h	42 ms
	06h	1.28 us	16h	84 ms
	07h	2.56 us	1Eh-17h	Reserved
	08h	5.12 us	1Fh	Reserved
	09h	10.2 us		
	0Ah	20.5 us		
	0Bh	41.0 us		
	0Ch	81.9 us		
	0Dh	163.8 us		
	0Eh	327.7 us		
	0Fh	655.4 us		
	Note:			
	1. This setting is not supported except as a DRAM scrub rate when no other memory accesses are being performed.			
23:5	Reserved.			
4:0	DramScrub: DRAM scrub rate. See: L3Scrub. BIOS: 2.10.5.6. Specifies time between 64 B scrub events. See D18F3x5C and D18F3x60.			

D18F3x5C DRAM Scrub Address Low

In addition to sequential DRAM scrubbing, the DRAM scrubber has a redirect mode for scrubbing DRAM locations accessed during normal operation. This is enabled by setting D18F3x5C[ScrubReDirEn]. When a DRAM read is generated by any agent other than the DRAM scrubber, correctable ECC errors are corrected as the data is passed to the requestor, but the data in DRAM is not corrected if redirect scrubbing mode is disabled. In scrubber redirect mode, correctable errors detected during normal DRAM read accesses redirect the scrubber to the location of the error. After the scrubber corrects the location in DRAM, it resumes scrubbing from where it left off. DRAM scrub address registers are not modified by the redirect scrubbing mode. Sequential scrubbing and scrubber redirection can be enabled independently or together. ECC errors detected by the scrubber are logged in the MCA registers (See D18F3x40 [MCA NB Control]).

Bits	Description
31:6	ScrubAddr[31:6]: DRAM scrubber address bits[31:6]. Read-write; updated-by-hardware. ScrubAddr[47:6] = {D18F3x60[ScrubAddr[47:32]], ScrubAddr[31:6]}. Reset: 0. ScrubAddr points to a DRAM cacheline in physical address space. BIOS should initialize the scrubber address register to the base address of the node specified by D18F1x[17C:140,7C:40] [DRAM Base/Limit] prior to enabling sequential scrubbing through D18F3x58[DramScrub]. When sequential scrubbing is enabled: it starts at the address that the scrubber address registers are initialized to; it increments through address space and updates the scrubber address registers as it does so; when the scrubber reaches the DRAM limit address specified by D18F1x[17C:140,7C:40], it wraps around to the base address. Reads of the scrubber address registers provide the next cacheline to be scrubbed.

5:1	Reserved.
0	ScrubReDirEn: DRAM scrubber redirect enable. Read-write. Reset: 0. BIOS: 2.10.5.6. If a correctable error is discovered from a non-scrubber DRAM read, then the data is corrected before it is returned to the requestor; however, the DRAM location may be left in a corrupted state (until the next time the scrubber address counts up to that location, if sequential scrubbing is enabled through D18F3x58 [DramScrub]). 1=Enables the scrubber to immediately scrub any address in which a correctable error is discovered. This bit and sequential scrubbing can be enabled independently or together; if both are enabled, the scrubber jumps from the scrubber address to where the correctable error was discovered, scrubs that location, and then jumps back to where it left off; the scrubber address register is not affected during scrubber redirection.

D18F3x60 DRAM Scrub Address High

Bits	Description
31:16	Reserved.
15:0	ScrubAddr[47:32]: DRAM scrubber address bits[47:32]. See: D18F3x5C [ScrubAddr[31:6]]. Reset: 0.

D18F3x64 Hardware Thermal Control (HTC)

See 2.11.2.1 [[PROCHOT_L and Hardware Thermal Control \(HTC\)](#)]. [D18F3x64](#) should only be programmed on internal node 0 ([D18F3xE8](#)[IntNodeNum]=00b) if [D18F3xE8](#)[MultiNodeCpu]=1. 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. BIOS: D18F3xDC [HwPstateMaxVal]. 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 [HwPstateMaxVal]. No P-state limit is applied if the value written is less than D18F4x15C [NumBoostStates]. See 2.11.2.1 [PROCHOT_L and Hardware Thermal Control (HTC)] and 2.5.2.1.2.2 [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 HtcTmpLmt minus HtcHystLmt. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0</td> </tr> <tr> <td>1h</td> <td>0.5</td> </tr> <tr> <td>Eh-2h</td> <td><HtcHystLmt*0.5></td> </tr> <tr> <td>Fh</td> <td>7.5</td> </tr> </tbody> </table>	Bits	Description	0h	0	1h	0.5	Eh-2h	<HtcHystLmt*0.5>	Fh	7.5
Bits	Description										
0h	0										
1h	0.5										
Eh-2h	<HtcHystLmt*0.5>										
Fh	7.5										
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="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>52</td> </tr> <tr> <td>01h</td> <td>52.5</td> </tr> <tr> <td>7Eh-02h</td> <td><(HtcTmpLmt*0.5) + 52></td> </tr> <tr> <td>7Fh</td> <td>115.5</td> </tr> </tbody> </table>	Bits	Description	00h	52	01h	52.5	7Eh-02h	<(HtcTmpLmt*0.5) + 52>	7Fh	115.5
Bits	Description										
00h	52										
01h	52.5										
7Eh-02h	<(HtcTmpLmt*0.5) + 52>										
7Fh	115.5										
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 MSRC001_0071[CurPstateLimit] changes due to SB-RMI, software P-state limit, or HTC. 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. It is cleared by writing a 1 to it.										
4	HtcAct: HTC-active state. Read-only. 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 2.11.2.2 [Software P-state Limit Control]. D18F3x68 should only be programmed on internal node 0 (D18F3xE8[IntNodeNum]=00b) if D18F3xE8[MultiNodeCpu]=1. 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 2.5.2.1.2.2 [Hardware P-state Numbering] and is 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

Read-write; Reset-applied.

- To ensure deadlock free operation the following minimum buffer allocations are required:
 - D18F3x6C[UpRspDBC] >= 1.

- $D18F3x6C[DnReqDBC] \geq 1$.
- $D18F3x6C[UpReqDBC] \geq 1$.
- $D18F3x6C[DnRspDBC] \geq 1$.
- If $D18F0x68[DispRefModeEn]$ is set or any of the $D18F0x[E4,C4,A4,84][IsocEn]$ bits are set: $IsocRspDBC \geq 1$.
- The total number of data buffers allocated in this register and $D18F3x7C$ must satisfy the following equation:
 - $D18F3x6C[UpReqDBC] + D18F3x6C[UpRspDBC] + D18F3x6C[DnReqDBC] + D18F3x6C[DnRspDBC] + D18F3x6C[IsocRspDBC] + (IF (D18F3x7C[Sri2XbarFreeRspDBC]==0) THEN (D18F3x7C[Sri2XbarFreeXreqDBC]*2) ELSE D18F3x7C[Sri2XbarFreeXreqDBC] ENDIF) + D18F3x7C[Sri2XbarFreeRspDBC] \leq 32$.

See 2.9.3.2.5.1 [Recommended Buffer Count Settings Overview].

Bits	Description
31	Reserved.
30:28	IsocRspDBC: isochronous response data buffer count. Cold reset: 3. BIOS: 1.
27:19	Reserved.
18:16	UpRspDBC: upstream response data buffer count. Cold reset: 2. BIOS: 1.
15:8	Reserved.
7:6	DnRspDBC: downstream response data buffer count. Cold reset: 2. BIOS: 1.
5:4	DnReqDBC: downstream request data buffer count. Cold reset: 1. BIOS: 1.
3	Reserved.
2:0	UpReqDBC: upstream request data buffer count. Cold reset: 2. BIOS: 2.

D18F3x70 SRI to XBAR Command Buffer Count

Read-write; Reset-applied.

- To ensure deadlock free operation the following minimum buffer allocations are required:
 - $D18F3x70[UpRspCBC] \geq 1$.
 - $D18F3x70[UpPreqCBC] \geq 1$.
 - $D18F3x70[DnPreqCBC] \geq 1$.
 - $D18F3x70[UpReqCBC] \geq 1$.
 - $D18F3x70[DnReqCBC] \geq 1$.
 - $D18F3x70[DnRspCBC] \geq 1$.
- If $D18F0x68[DispRefModeEn]$ is set or any of the $D18F0x[E4,C4,A4,84][IsocEn]$ bits are set: $IsocReqCBC \geq 1$ $IsocRspCBC \geq 1$
- If any of the $D18F0x[E4,C4,A4,84][IsocEn]$ bits are set and isochronous posted requests may be generated by the system: $IsocPreqCBC \geq 1$
- The total number of SRI to XBAR command buffers allocated in this register and $D18F3x7C$ must satisfy the following equation:
 - $D18F3x70[IsocRspCBC] + D18F3x70[IsocPreqCBC] + D18F3x70[IsocReqCBC] + D18F3x70[UpRspCBC] + D18F3x70[DnPreqCBC] + D18F3x70[UpPreqCBC] + D18F3x70[DnReqCBC] + D18F3x70[DnRspCBC] + D18F3x70[UpReqCBC] + IF (D18F3x7C[Sri2XbarFreeRspCBC]==0) THEN (2*D18F3x7C[Sri2XbarFreeXreqCBC]) ELSE (D18F3x7C[Sri2XbarFreeRspCBC] + D18F3x7C[Sri2XbarFreeXreqCBC]) ENDIF) \leq 48$.

See 2.9.3.2.5.1 [Recommended Buffer Count Settings Overview].

Bits	Description
31	Reserved.
30:28	IsocRspCBC: isoc response command buffer count. Cold reset: 4. BIOS: 1.
27	Reserved.
26:24	IsocPreqCBC: isoc posted request command buffer count. Cold reset: 1. BIOS: 0.
23	Reserved.
22:20	IsocReqCBC: isoc request command buffer count. Cold reset: 5. BIOS: 1.
19	Reserved.
18:16	UpRspCBC: upstream response command buffer count. Cold reset: 3. BIOS: 7.
15	Reserved.
14:12	DnPreqCBC: downstream posted request command buffer count. Cold reset: 3. BIOS: 1.
11	Reserved.
10:8	UpPreqCBC: upstream posted request command buffer count. Cold reset: 3. BIOS: 1.
7:6	DnRspCBC: downstream response command buffer count. Cold reset: 2. BIOS: 1.
5:4	DnReqCBC: downstream request command buffer count. Cold reset: 2. BIOS: 1.
3	Reserved.
2:0	UpReqCBC: upstream request command buffer count. Cold reset: 3. BIOS: 5.

D18F3x74 XBAR to SRI Command Buffer Count

Read-write; Reset-applied.

Table 218: XBAR Definitions

Term	Definition
SpqSize	Probe command queue size. SpqSize = 16.
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 in a multinode system the following minimum buffer allocations are required:
 - ProbeCBC ≥ 2 .
 - DnPreqCBC ≥ 1 if D18F0x60[NodeCnt] specifies 1 node; DnPreqCBC ≥ 2 if D18F0x60[NodeCnt] specifies ≥ 2 nodes.
 - UpPreqCBC ≥ 1 .
 - DnReqCBC ≥ 1 .
 - UpReqCBC ≥ 1 .
- To ensure deadlock free operation in a single-node system the following minimum buffer allocations are required:
 - ProbeCBC ≥ 2
 - UpReqCBC ≥ 1
 - UpPreqCBC ≥ 1
- If D18F0x68[DispRefModeEn] is set or the node is directly connected to an IO link with an IOMMU present:
 - IsocReqCBC ≥ 1

- If any of the **D18F0x[E4,C4,A4,84][IsocEn]** bits are set:
IsocReqCBC >= 1
- If any of the **D18F0x[E4,C4,A4,84][IsocEn]** bits are set and isochronous posted requests may be generated by the system:
IsocPreqCBC >= 1
- The total number of SRQ (XBAR to SRI command) buffers allocated:
 - $(\text{D18F3x1A0}[\text{L3ToSriReqCBC}] + \text{PrbRsp} + \text{D18F3x74}[\text{UpReqCBC}] + \text{D18F3x74}[\text{UpPreqCBC}] + \text{D18F3x74}[\text{DnReqCBC}] + \text{D18F3x74}[\text{DnPreqCBC}] + \text{D18F3x74}[\text{IsocReqCBC}] + \text{D18F3x74}[\text{IsocPreqCBC}] + \text{D18F3x74}[\text{DRReqCBC}] + \text{D18F3x7C}[\text{Xbar2SriFreeListCBC}] + \text{D18F3x7C}[\text{SrqExtFreeListBC}]) \leq \text{SrqSize}.$
 - $((\text{D18F3x1A0}[\text{CpuCmdBufCnt}] * \text{NumOfCompUnitsOnNode}) + \text{PrbRsp} + \text{D18F3x74}[\text{UpReqCBC}] + \text{D18F3x74}[\text{UpPreqCBC}] + \text{D18F3x74}[\text{DnReqCBC}] + \text{D18F3x74}[\text{DnPreqCBC}] + \text{D18F3x74}[\text{IsocReqCBC}] + \text{D18F3x74}[\text{IsocPreqCBC}] + \text{D18F3x74}[\text{DRReqCBC}] + \text{D18F3x7C}[\text{Xbar2SriFreeListCBC}] + \text{D18F3x7C}[\text{SrqExtFreeListBC}]) \leq \text{SrqSize}.$
- The total number of SPQ (probe command) buffers allocated must satisfy the following equation:
 - $(\text{D18F3x17C}[\text{SPQPrbFreeCBC}] + \text{D18F3x74}[\text{ProbeCBC}]) \leq \text{SpqSize}.$

See 2.9.3.2.5.1 [Recommended Buffer Count Settings Overview].

Bits	Description										
31:28	DRReqCBC: display refresh request command buffer count. Cold reset: 0. BIOS: 0.										
27	Reserved.										
26:24	IsocPreqCBC: isochronous posted request command buffer count. Cold reset: 0. BIOS: 0.										
23:20	IsocReqCBC: isochronous request command buffer count. Cold reset: 4. BIOS: 1.										
19:16	ProbeCBC: probe command buffer count. Cold reset: 7. BIOS: 7. <table border="1" style="margin-left: 20px;"> <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>Dh</td> <td>13 buffers</td> </tr> <tr> <td>Fh-Eh</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	0h	0 buffers	Ch-1h	<ProbeCBC> buffers	Dh	13 buffers	Fh-Eh	Reserved.
Bits	Description										
0h	0 buffers										
Ch-1h	<ProbeCBC> buffers										
Dh	13 buffers										
Fh-Eh	Reserved.										
15	Reserved.										
14:12	DnPreqCBC: downstream posted request command buffer count. Cold reset: 1. BIOS: 2.										
11	Reserved.										
10:8	UpPreqCBC: upstream posted request command buffer count. Cold reset: 1. BIOS: 1.										
7	Reserved.										
6:4	DnReqCBC: downstream request command buffer count. Cold reset: 1. BIOS: 1.										
3	Reserved.										
2:0	UpReqCBC: upstream request command buffer count. Cold reset: 1. BIOS: 1.										

D18F3x78 MCT to XBAR Buffer Count

Read-write; Reset-applied.

- To ensure deadlock free operation the following minimum buffer allocations are required:
 ProbeCBC >= 1 RspCBC >= 1 RspDBC >= 2
 RspDBC >= **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 >= Dh

- The total number of command buffers allocated in this register must satisfy the following equation:
 - $(D18F3x78[ProbeCBC] + D18F3x78[RspCBC]) \leq 32$.

Bits	Description												
31:22	Reserved.												
21:16	RspDBC: response data buffer count. 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. Cold reset: Ch. BIOS: IF (PrbFltrEn) THEN 0Ch ELSE 0Eh ENDIF.												
7:5	Reserved.												
4:0	RspCBC: response command buffer count. Cold reset: 14h. BIOS: IF (PrbFltrEn) THEN 14h ELSE 12h ENDIF.												

D18F3x7C Free List Buffer Count

Reset-applied.

- 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] >= 2.

See 2.9.3.2.5.1 [Recommended Buffer Count Settings Overview].

Bits	Description
31	Reserved.
30:28	Xbar2SriFreeListCBInc: XBAR to SRI free list command buffer increment. Read-write. Cold reset: 0. This is use to add buffers to the free list pool if they are reclaimed from hard allocated entries without having to go through warm reset.
27	Reserved.
26:23	SrqExtFreeListBC: extend SRQ freelist tokens. Read-write. Cold reset: 8h. BIOS: 8h. Can only be used by requests from cores to L3 or DRAM.
22:20	Sri2XbarFreeRspDBC: SRI to XBAR free response data buffer count. Read-write. Cold reset: 3. BIOS: 0.
19:16	Sri2XbarFreeXreqDBC: SRI to XBAR free request and posted request data buffer count. Read-write. Cold reset: 3h. BIOS: Dh. If (Sri2XbarFreeRspDBC==0) 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: 0h.
11:8	Sri2XbarFreeXreqCBC: SRI to XBAR free request and posted request command buffer count. Read-write. Cold reset: Bh. BIOS: Fh. (30 buffers) If (Sri2XbarFreeRspCBC==0) 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: 14h. BIOS: 16h.

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. If an NB COF-change command is issued (NbCofChg), the NB returns in the new NB P-state.

In multi-node systems, these registers should be programmed identically in all nodes.

Table 219: 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. • See 2.9.5 [Memory Scrubbers] for recommended clock divisor settings when memory scrubbing is enabled.
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 2.5.2.2 [NB P-states].

Table 219: SMAF Action Definition

Register	SmafAct	ACPI state	Description
D18F3x80[15:8]	SmafAct1	C1E, or Link init.	Initiated by an access to the ACPI-defined P_LVL3 register or in response to a write to the Link Frequency Change and Resize LDTSTOP_L Command register in the IO hub. LDTSTOP_L is expected to be asserted while in this state. See also 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E].
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. BIOS: 111b.
28	Reserved.
27	NbCofChgSmafAct3. See: NbCofChgSmafAct0. BIOS: 0.
26	NbGateEnSmafAct3. See: NbGateEnSmafAct0. BIOS: 0.
25	NbLowPwrEnSmafAct3. See: NbLowPwrEnSmafAct0. BIOS: 1.
24	CpuPrbEnSmafAct3. See: CpuPrbEnSmafAct0. BIOS: 0.
23:21	ClkDivisorSmafAct2. See: ClkDivisorSmafAct0. BIOS: 0.
20	Reserved.
19	NbCofChgSmafAct2. See: NbCofChgSmafAct0. BIOS: 1.
18	NbGateEnSmafAct2. See: NbGateEnSmafAct0. BIOS: 0.
17	NbLowPwrEnSmafAct2. See: NbLowPwrEnSmafAct0. BIOS: 1.
16	CpuPrbEnSmafAct2. See: CpuPrbEnSmafAct0. BIOS: 1.
15:13	ClkDivisorSmafAct1. See: ClkDivisorSmafAct0. BIOS: 111b.
12	Reserved.
11	NbCofChgSmafAct1. See: NbCofChgSmafAct0. BIOS: 0.
10	NbGateEnSmafAct1. See: NbGateEnSmafAct0. BIOS: 0.
9	NbLowPwrEnSmafAct1. See: NbLowPwrEnSmafAct0. BIOS: 1.
8	CpuPrbEnSmafAct1. See: CpuPrbEnSmafAct0. BIOS: 0.

7:5	<p>ClkDivisorSmafAct0: clock divisor. 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" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">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>Divide-by 128</td> </tr> <tr> <td>010b</td> <td>Divide-by 4</td> <td>110b</td> <td>Divide-by 512</td> </tr> <tr> <td>011b</td> <td>Divide-by 8</td> <td>111b</td> <td>Turn off clocks</td> </tr> </tbody> </table> <p>See 2.9.5 [Memory Scrubbers].</p>	Bits	Description	Bits	Description	000b	Divide-by 1	100b	Divide-by 16	001b	Divide-by 2	101b	Divide-by 128	010b	Divide-by 4	110b	Divide-by 512	011b	Divide-by 8	111b	Turn off clocks
Bits	Description	Bits	Description																		
000b	Divide-by 1	100b	Divide-by 16																		
001b	Divide-by 2	101b	Divide-by 128																		
010b	Divide-by 4	110b	Divide-by 512																		
011b	Divide-by 8	111b	Turn off clocks																		
4	Reserved.																				
3	<p>NbCofChgSmafAct0: Northbridge FID change. 1=The NB P-state is changed after LDTSTOP_L is asserted. If this bit is set, then NbLowPwrEn must be set and ClkDivisor must be 000b. See 2.5.2.2 [NB P-states].</p>																				
2	<p>NbGateEnSmafAct0: northbridge gate enable. This bit does not control hardware. NbLowPwrEn is required to be set if this bit is set.</p>																				
1	<p>NbLowPwrEnSmafAct0: Northbridge low-power enable. 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.</p>																				
0	<p>CpuPrbEnSmafAct0: CPU direct probe enable. 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 and ClkDivisorSmafAct0!=000b, 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 or ClkDivisorSmafAct0=000b, 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]. BIOS: 000b.
28	Reserved.
27	NbCofChgSmafAct7. See: D18F3x80[NbCofChgSmafAct0]. BIOS: 0.
26	NbGateEnSmafAct7. See: D18F3x80[NbGateEnSmafAct0]. BIOS: 0.

25	NbLowPwrEnSmafAct7. See: D18F3x80[NbLowPwrEnSmafAct0] . BIOS: 0.
24	CpuPrbEnSmafAct7. See: D18F3x80[CpuPrbEnSmafAct0] . BIOS: 1.
23:21	ClkDivisorSmafAct6. See: D18F3x80[ClkDivisorSmafAct0] . BIOS: 111b.
20	Reserved.
19	NbCofChgSmafAct6. See: D18F3x80[NbCofChgSmafAct0] . BIOS: 0.
18	NbGateEnSmafAct6. See: D18F3x80[NbGateEnSmafAct0] . BIOS: 0.
17	NbLowPwrEnSmafAct6. See: D18F3x80[NbLowPwrEnSmafAct0] . BIOS: 1.
16	CpuPrbEnSmafAct6. See: D18F3x80[CpuPrbEnSmafAct0] . BIOS: 0.
15:8	Reserved.
7:5	ClkDivisorSmafAct4. See: D18F3x80[ClkDivisorSmafAct0] . BIOS: 111b.
4	Reserved.
3	NbCofChgSmafAct4. See: D18F3x80[NbCofChgSmafAct0] . BIOS: 0.
2	NbGateEnSmafAct4. See: D18F3x80[NbGateEnSmafAct0] . BIOS: 0.
1	NbLowPwrEnSmafAct4. See: D18F3x80[NbLowPwrEnSmafAct0] . BIOS: 1.
0	CpuPrbEnSmafAct4. See: D18F3x80[CpuPrbEnSmafAct0] . BIOS: 0.

D18F3x88 NB Configuration 1 Low (NB_CFG1_LO)

Bits	Description
31:0	See: MSRC001_001F[31:0] .

D18F3x8C NB Configuration 1 High (NB_CFG1_HI)

Bits	Description
31:0	See: MSRC001_001F[63:32] .

D18F3x90 GART Aperture Control

Read-write; **Same-for-all**. Reset: 0000_0000h. See [2.9.2 \[GART\]](#).

Bits	Description
31:7	Unused.
6	DisGartTblWlkPrb: disable GART table walk probes. 1=Disables generation of probes for GART table walks.
5	DisGartIo: disable GART IO accesses. 1=Disables requests from IO devices from accessing the GART.
4	DisGartCpu: disable GART CPU accesses. 1=Disables requests from CPUs from accessing the GART.

3:1	GartSize: GART size. Specifies the size of address space allocated to the GART.			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	000b	32 MB	100b	512 MB
	001b	64 MB	101b	1 GB
	010b	128 MB	110b	2 GB
	011b	256 MB	111b	Reserved
0	GartEn: GART enable. 1=Enables GART address translation for accesses falling within the GART aperture. D18F3x94 [GartAperBaseAddr] and other related registers should be initialized before GartEn is set.			

D18F3x94 GART Aperture Base

Read-write; Same-for-all. See 2.9.2 [GART].

Bits	Description
31:15	Reserved.
14:0	GartAperBaseAddr[39:25]: GART aperture base address bits[39:25]. Reset: X. Specifies the base address of the GART aperture range. Based on D18F3x90 [GartSize], some of the LSB address bits are assumed to be 0 (e.g., if the GART is 1 Gbyte, then only GartAperBaseAddr[39:30] is meaningful). This field along with D18F3x90 [GartSize] specifies the GART aperture address range. BIOS can place the GART aperture below the 4-gigabyte level in address space in order to support legacy operating systems and legacy AGP cards (that do not support 64-bit address space). Note: GART apertures above 1 terabyte are not supported.

D18F3x98 GART Table Base

Read-write; Same-for-all. See 2.9.2 [GART].

Bits	Description												
31:4	GartTblBaseAddr[39:12]: GART table base address bits[39:12]. Reset: X. Specifies the base address of the table of GART page table entries (PTEs) used in GART address translation. Accesses to the GART aperture address range specified by D18F3x90 and D18F3x94 , address GA[39:0], are translated to the physical address specified by the corresponding GART PTE. Each PTE is 32-bits wide. The first PTE corresponds to the first 4 Kbyte page of the GART aperture, and so on. PTEs are defined as follows: <table border="1" style="margin-left: 20px;"> <tr> <th>PTE bits</th> <th>Description</th> </tr> <tr> <td>31:12</td> <td>Physical address bits[31:12]</td> </tr> <tr> <td>11:4</td> <td>Physical address bits[39:32]</td> </tr> <tr> <td>3:2</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>Coherent: 1=Probes are required for accesses to the range.</td> </tr> <tr> <td>0</td> <td>Valid: 1=Entry is valid.</td> </tr> </table>	PTE bits	Description	31:12	Physical address bits[31:12]	11:4	Physical address bits[39:32]	3:2	Reserved	1	Coherent: 1=Probes are required for accesses to the range.	0	Valid: 1=Entry is valid.
PTE bits	Description												
31:12	Physical address bits[31:12]												
11:4	Physical address bits[39:32]												
3:2	Reserved												
1	Coherent: 1=Probes are required for accesses to the range.												
0	Valid: 1=Entry is valid.												
3:0	Reserved.												

D18F3x9C GART Cache Control

Same-for-all. Reset: 0000_0000h. See 2.9.2 [GART].

Bits	Description
31:2	Reserved.

1	GartPteErr: GART PTE error. Read; set-by-hardware; write-1-to-clear. 1=An invalid PTE was encountered during a table walk.
0	InvGart: invalidate GART. Read-write; cleared-by-hardware. Setting this bit causes the GART cache to be invalidated. This bit is cleared by hardware when the invalidation is complete. Software writing this bit to 0 has no effect.

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 00h=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 2.5.2.1.8 [BIOS Requirements for Core P-state Initialization and Transitions] .																				
30:29	Reserved.																				
28	NbPstateForce: NB P-state force on next LDTSTOP_L assertion. Read-write. Reset: 0. 1=When there is a pending NB P-state change (after the FIDVID broadcast), the NB P-state is changed on the next LDTSTOP_L assertion regardless of whether D18F3x[84:80] indexed by the SMAF code corresponding to the assertion indicates an NB P-state change.																				
27:16	ConfigId: Configuration identifier. Read-only. Reset: Product-specific. Specifies the configuration ID associated with the product.																				
15:14	Reserved.																				
13:11	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. <table border="1" data-bbox="267 1333 1047 1507"> <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	IdleExitEn: idle exit enable. Read-write. Reset: 0. BIOS: 1. 1=Enable BP[5] to function as the IDLE_EXIT_L pin. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E] .																				
9	SviHighFreqSel: SVI high frequency select. Read-write. Cold reset: 0. 0=400 KHz. 1=3.4 MHz. Writes to this field take effect at the next SVI command boundary. If 3.4 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.																				
8	Reserved.																				

7	PsiVidEn: PSI_L VID enable. Read-write. Reset: 0. 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 2.5.1.4.1 [PSI_L Bit].
6:0	PsiVid: PSI_L VID threshold. Read-write. Reset: 0. When enabled by PsiVidEn, this field specifies the threshold value of VID code generated by the processor, which in turn determines the state of PSI_L. When the VID code generated by the processor is less than PsiVid (i.e., the VID code is specifying a higher voltage level than the PsiVid-specified voltage level), then PSI_L is high; when the VID code is greater than or equal to PsiVid, PSI_L is driven low. See 2.5.1.4.1 [PSI_L Bit].

D18F3xA4 Reported Temperature Control

The slew rate controls in this register are used to filter processor temperature measurements. Separate controls are provided for a measured temperature that is higher or lower than Tctl. The per-step timer counts as long as the measured temperature stays either above or below Tctl. Each time the measured temperature changes to the other side of Tctl, the step timer resets, and Tctl is not changed. If, for example, step times are enabled in both directions, Tctl=62.625, and the measured temperature keeps jumping quickly between 62.5 and 63.0, then (assuming the step times are long enough) Tctl would not change. However, once the measured temperature settles on one side of Tctl, Tctl can step toward the measured temperature. If the difference of measured temperature minus Tctl is greater than the value set by MaxTmpDiffUp, then Tctl is set equal to the measured temperature. See 2.11 [Thermal Functions].

Bits	Description																				
31:21	<p>CurTmp: current temperature. IF (D18F3xA4[CurTmpTjSel]==11b) THEN Read-write. ELSE Read-only. ENDIF. Reset: X. Provides the current control temperature, Tctl (after the slew-rate controls have been applied). See 2.11.1 [The Tctl Temperature Scale]. See CurTmpTjSel. If D18F3xE8[MultiNodeCpu]=1, this field should be accessed on internal node 0 (D18F3xE8[IntNodeNum]=00b) to obtain temperature information for the processor. Reading this field on internal node 1 provides only the temperature for internal node 1. If CurTmpTjSel=11b, programming this field on internal node 1 does not affect CurTmp or temperature-driven logic on internal node 0.</p> <p>IF (D18F3xA4[CurTmpTjSel]!=11b) THEN</p> <table border="0"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></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="0"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000h-003h</td> <td>-49</td> </tr> <tr> <td>004h-007h</td> <td>-48.5</td> </tr> <tr> <td>7FBh-008h</td> <td><(CurTmp[10:2]*0.5)-49></td> </tr> <tr> <td>7FFh-7FCh</td> <td>206.5</td> </tr> </tbody> </table> <p>ENDIF.</p>	<u>Bits</u>	<u>Description</u>	000h	0	001h	0.125	7FEh-002h	<CurTmp*0.125>	7FFh	255.875	<u>Bits</u>	<u>Description</u>	000h-003h	-49	004h-007h	-48.5	7FBh-008h	<(CurTmp[10:2]*0.5)-49>	7FFh-7FCh	206.5
<u>Bits</u>	<u>Description</u>																				
000h	0																				
001h	0.125																				
7FEh-002h	<CurTmp*0.125>																				
7FFh	255.875																				
<u>Bits</u>	<u>Description</u>																				
000h-003h	-49																				
004h-007h	-48.5																				
7FBh-008h	<(CurTmp[10:2]*0.5)-49>																				
7FFh-7FCh	206.5																				
20:18	Reserved.																				

17:16	<p>CurTmpTjSel: Current temperature select. Read-write. Reset: 00. These bits may be used for diagnostic software.</p> <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 LSBs 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 LSBs 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 LSBs are read-only zero.										
15:13	Reserved.										
12:8	<p>PerStepTimeDn[4:0]: per 1/8th step time down. Read-write. Cold reset: 18h. BIOS: 0Fh. Specifies the time per 1/8th step of Tctl when the measured temperature is less than the Tctl. It is encoded the same as PerStepTimeUp.</p>										
7	<p>TmpSlewDnEn: temperature slew downward enable. Read-write. Cold reset: 0. BIOS: 1. 1=Slew rate controls in the downward direction are enabled. 0=Downward slewing disabled; if the measured temperature is detected to be less than Tctl then Tctl is updated to match the measured temperature.</p>										
6:5	<p>TmpMaxDiffUp: temperature maximum difference up. Read-write. Cold reset: 00b. BIOS: 11b. Specifies the maximum difference between Tctl and the measured temperature, when the measured value is greater than Tctl (i.e., when the temperature has risen). If this difference exceeds the specified value, Tctl jumps to the measured temperature value.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Upward slewing disabled; if the measured temperature is detected to be greater than Tctl then Tctl is updated to match the measured temperature.</td> </tr> <tr> <td>01b</td> <td>Tctl is held to less than or equal to measured temperature minus 1.0.</td> </tr> <tr> <td>10b</td> <td>Tctl is held to less than or equal to measured temperature minus 3.0.</td> </tr> <tr> <td>11b</td> <td>Tctl is held to less than or equal to measured temperature minus 9.0.</td> </tr> </tbody> </table>	Bits	Description	00b	Upward slewing disabled; if the measured temperature is detected to be greater than Tctl then Tctl is updated to match the measured temperature.	01b	Tctl is held to less than or equal to measured temperature minus 1.0.	10b	Tctl is held to less than or equal to measured temperature minus 3.0.	11b	Tctl is held to less than or equal to measured temperature minus 9.0.
Bits	Description										
00b	Upward slewing disabled; if the measured temperature is detected to be greater than Tctl then Tctl is updated to match the measured temperature.										
01b	Tctl is held to less than or equal to measured temperature minus 1.0.										
10b	Tctl is held to less than or equal to measured temperature minus 3.0.										
11b	Tctl is held to less than or equal to measured temperature minus 9.0.										
4:0	<p>PerStepTimeUp[4:0]: per 1/8th step time up. Read-write. Cold reset: 00h. BIOS: 0Fh. Specifies the time per 1/8th step of Tctl when the measured temperature is greater than the reported temperature.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>07h-00h</td> <td><PerStepTimeUp[2:0] * 0.5> ms, ranging from 0.5 to 4 ms.</td> </tr> <tr> <td>0Fh-08h</td> <td><PerStepTimeUp[2:0] * 5> ms, ranging from 5 to 40 ms.</td> </tr> <tr> <td>17h-10h</td> <td><PerStepTimeUp[2:0] * 50> ms, ranging from 50 to 400 ms.</td> </tr> <tr> <td>1Fh-18h</td> <td><PerStepTimeUp[2:0] * 0.5> s, ranging from 0.5 to 4 s.</td> </tr> </tbody> </table>	Bits	Description	07h-00h	<PerStepTimeUp[2:0] * 0.5> ms, ranging from 0.5 to 4 ms.	0Fh-08h	<PerStepTimeUp[2:0] * 5> ms, ranging from 5 to 40 ms.	17h-10h	<PerStepTimeUp[2:0] * 50> ms, ranging from 50 to 400 ms.	1Fh-18h	<PerStepTimeUp[2:0] * 0.5> s, ranging from 0.5 to 4 s.
Bits	Description										
07h-00h	<PerStepTimeUp[2:0] * 0.5> ms, ranging from 0.5 to 4 ms.										
0Fh-08h	<PerStepTimeUp[2:0] * 5> ms, ranging from 5 to 40 ms.										
17h-10h	<PerStepTimeUp[2:0] * 50> ms, ranging from 50 to 400 ms.										
1Fh-18h	<PerStepTimeUp[2:0] * 0.5> s, ranging from 0.5 to 4 s.										

D18F3xA8 Pop Up and Down P-states

Bits	Description
31:29	<p>PopDownPstate. Read-write. Reset: D18F3xDC[HwPstateMaxVal]. BIOS: D18F3xDC[HwPstateMaxVal]. Specifies the pop-down P-state number. This field uses hardware P-state numbering. See 2.5.3.3.3 [Core C6 (CC6) State] .</p>
28:0	Reserved.

D18F3xB0 On-Line Spare Control

Bits	Description
31:28	Reserved.

27:24	EccErrCnt: ECC error count. IF (EccErrCntWrEn) THEN Read-write. ELSE Read-only. ENDIF. Reset: 0. Specifies the number of ECC errors for the chip select selected by EccErrCntDramCs and EccErrCntChan. This field can be written by software to clear the count. This field returns Fh if 15 or more correctable ECC errors have occurred.										
23	EccErrCntWrEn: ECC error counter write enable. Read-write. Reset: 0. 1=Enable writes to the EccErrCnt field.										
22:21	Reserved.										
20	EccErrCntDct: ECC error counter DCT. Read-write. Reset: 0. Specifies the DCT for which ECC error count information is returned in the EccErrCnt field. 0=DCT 0. 1=DCT 1.										
19:16	EccErrCntDramCs: ECC error counter DRAM chip select. Read-write. Reset: 0. Specifies the DRAM chip select for which ECC error count information is returned in the EccErrCnt field. See D18F2x[5C:40]_dct[1:0] [DRAM CS Base Address] . <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>CS0</td> </tr> <tr> <td>6h-1h</td> <td>CS<EccErrCntDramCs></td> </tr> <tr> <td>7h</td> <td>CS7</td> </tr> <tr> <td>Fh-8h</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	CS0	6h-1h	CS<EccErrCntDramCs>	7h	CS7	Fh-8h	Reserved
<u>Bits</u>	<u>Description</u>										
0h	CS0										
6h-1h	CS<EccErrCntDramCs>										
7h	CS7										
Fh-8h	Reserved										
15:14	EccErrInt: ECC error interrupt type. Read-write. Reset: 0. Specifies the type of interrupt generated when the EccErrCnt field for any chip select and channel transitions to 1111b. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No Interrupt.</td> </tr> <tr> <td>01b</td> <td>Reserved</td> </tr> <tr> <td>10b</td> <td>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>	<u>Bits</u>	<u>Description</u>	00b	No Interrupt.	01b	Reserved	10b	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
<u>Bits</u>	<u>Description</u>										
00b	No Interrupt.										
01b	Reserved										
10b	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										
13:12	SwapDoneInt: swap complete interrupt type. Read-write. Reset: 0. Specifies the type of interrupt generated when a swap is complete. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No Interrupt.</td> </tr> <tr> <td>01b</td> <td>Reserved</td> </tr> <tr> <td>10b</td> <td>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>	<u>Bits</u>	<u>Description</u>	00b	No Interrupt.	01b	Reserved	10b	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
<u>Bits</u>	<u>Description</u>										
00b	No Interrupt.										
01b	Reserved										
10b	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										
11	Reserved.										
10:8	BadDramCs1: DCT1 bad DRAM chip select. See: BadDramCs0.										
7	Reserved.										
6:4	BadDramCs0: DCT0 bad DRAM chip select. Read-write. Cold reset: 0. Programmed with the DRAM chip select to be replaced when SwapEn is set. Cannot be written when SwapDone is set.										
3	SwapDone1: DCT1 swap complete. See: SwapDone0.										
2	SwapEn1: DCT1 swap enable. See: SwapEn0.										

1	SwapDone0: DCT0 swap complete. Read-write; set-by-hardware. Cold reset: 0. 1=The hardware has completed copying the data to the spare rank. This bit can also be set by BIOS to immediately enable the swap to the spare rank after suspend to RAM. Once this bit is set it cannot be cleared by software. This bit cannot be set by software if DRAM is enabled D18F2x110 [DramEnable].
0	SwapEn0: DCT0 swap enable. Read; write-1-only. Reset: 0. Setting this bit causes the hardware to copy the contents of the DRAM chip select identified by BadDramCs to the spare rank. The DRAM scrubber (D18F3x5C , D18F3x60) must be enabled with a scrub address range that encompasses the address of the bad chip select for the swap to occur. The scrub rate is accelerated automatically by hardware until the copy completes, at which point the scrub rate returns to normal. During the copy, DRAM accesses (including accesses to the bad CS) proceed normally. Once this bit is set, it cannot be cleared by software.

D18F3xB8 NB Array Address

Reset: xxxx_xxxxh. [D18F3xB8](#) [NB Array Address] and [D18F3xBC](#) [NB Array Data Port] provide a mechanism to inject errors into DRAM and data read from internal NB arrays.

[D18F3xB8](#) should first be written with the target array and address within the array. Read and write accesses to [D18F3xBC](#) then access the target address within the target array.

Bits	Description																		
31:28	ArraySelect. Read-write. Selects the NB array to access. <table border="0"> <tr> <td><u>Bits</u></td> <td><u>Array</u></td> </tr> <tr> <td>7h-0h</td> <td>Reserved</td> </tr> <tr> <td>8h</td> <td>DRAM ECC [D18F3xBC_x8]</td> </tr> <tr> <td>Fh-9h</td> <td>Reserved</td> </tr> </table>	<u>Bits</u>	<u>Array</u>	7h-0h	Reserved	8h	DRAM ECC [D18F3xBC_x8]	Fh-9h	Reserved										
<u>Bits</u>	<u>Array</u>																		
7h-0h	Reserved																		
8h	DRAM ECC [D18F3xBC_x8]																		
Fh-9h	Reserved																		
27:10	Reserved.																		
9:0	ArrayAddress. Read-write. Selects the location to access within the selected array. <table border="0"> <tr> <td><u>ArraySelect</u></td> <td><u>Description</u></td> </tr> <tr> <td>8h</td> <td>DRAM ECC</td> </tr> <tr> <td></td> <td><u>Bit</u> <u>Description</u></td> </tr> <tr> <td></td> <td>[9:3] Reserved</td> </tr> <tr> <td></td> <td>[2:1] Select 16 byte quadrant in 64 byte cache line.</td> </tr> <tr> <td></td> <td>[0] Reserved</td> </tr> <tr> <td>All others:</td> <td></td> </tr> <tr> <td></td> <td><u>Bit</u> <u>Description</u></td> </tr> <tr> <td></td> <td>[9:0] Reserved</td> </tr> </table>	<u>ArraySelect</u>	<u>Description</u>	8h	DRAM ECC		<u>Bit</u> <u>Description</u>		[9:3] Reserved		[2:1] Select 16 byte quadrant in 64 byte cache line.		[0] Reserved	All others:			<u>Bit</u> <u>Description</u>		[9:0] Reserved
<u>ArraySelect</u>	<u>Description</u>																		
8h	DRAM ECC																		
	<u>Bit</u> <u>Description</u>																		
	[9:3] Reserved																		
	[2:1] Select 16 byte quadrant in 64 byte cache line.																		
	[0] Reserved																		
All others:																			
	<u>Bit</u> <u>Description</u>																		
	[9:0] Reserved																		

D18F3xBC NB Array Data Port

See [D18F3xB8](#) for register access information.

D18F3xBC_x8 DRAM ECC

This register controls injection of errors in writes to DRAM. See [2.13.3.1 \[DRAM Error Injection\]](#).

Bits	Description																				
31:29	Reserved.																				
28:20	ErrInjEn: enable error injection to word. Read-write. Reset: 0. Each bit in this field corresponds to a 16-bit DRAM word and enables injecting errors in that word. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Data[15:0]</td> </tr> <tr> <td>[1]</td> <td>Data[31:16]</td> </tr> <tr> <td>[2]</td> <td>Data[47:32]</td> </tr> <tr> <td>[3]</td> <td>Data[63:48]</td> </tr> <tr> <td>[4]</td> <td>Data[79:64]</td> </tr> <tr> <td>[5]</td> <td>Data[95:80]</td> </tr> <tr> <td>[6]</td> <td>Data[111:96]</td> </tr> <tr> <td>[7]</td> <td>Data[127:112]</td> </tr> <tr> <td>[8]</td> <td>ECC[15:0]</td> </tr> </tbody> </table>	Bit	Description	[0]	Data[15:0]	[1]	Data[31:16]	[2]	Data[47:32]	[3]	Data[63:48]	[4]	Data[79:64]	[5]	Data[95:80]	[6]	Data[111:96]	[7]	Data[127:112]	[8]	ECC[15:0]
Bit	Description																				
[0]	Data[15:0]																				
[1]	Data[31:16]																				
[2]	Data[47:32]																				
[3]	Data[63:48]																				
[4]	Data[79:64]																				
[5]	Data[95:80]																				
[6]	Data[111:96]																				
[7]	Data[127:112]																				
[8]	ECC[15:0]																				
19	Reserved.																				
18	DramErrEn. Read-write. Reset: 0. 1=Errors are continually injected on DRAM writes. The error injection takes place only on DRAM write accesses and should be initiated by a non-cacheable store. Errors continue to be injected on writes until this bit is cleared to a 0 by software.																				
17	EccWrReq. Read; write-1-only; cleared-by-hardware. Reset: 0. 1=Error is injected on DRAM write at the bits enabled by ErrInjEn and EccVector. A single error injection takes place on the next DRAM write access and should be initiated by a non-cacheable store. This bit is cleared by hardware after the write.																				
16	EccRdReq. Read; write-1-only; cleared-by-hardware. Reset: 0. 1=Indicates a DRAM ECC read is requested. The read takes place on the next DRAM read access and should be initiated by a non-cacheable load. The ECC bits read from DRAM are stored in EccVector. This bit is cleared by hardware after the read.																				
15:0	EccVector: error injection vector. Read-write. Reset: x. When used in conjunction with EccWrReq, each bit of EccVector enables injecting errors to the corresponding bit within each word enabled by ErrInjEn. When used in conjunction with EccRdReq, EccVector holds the contents of the DRAM ECC bits after the read.																				

D18F3xC4 SBI P-state Limit

See [MSRC001_0072 \[SBI P-state Limit\]](#).

Bits	Description
31:11	Reserved.
10:8	PstateLimit: P-state limit select. Read-only. See: MSRC001_0072[PstateLimit] . Uses hardware P-state numbering. See MSRC001_0072[PstateLimit] .
7:1	Reserved.
0	PstateLimitEn: P-state limit enable. Read-only. See: MSRC001_0072[PstateLimitEn] . 1=PstateLimit is enabled.

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_L 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"> • 200 MHz * (4 + D18F5x1[6C:60][NbFid]). <p>If D18F5x1[6C:60][NbDid] of the current NB 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>Reserved.</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	Reserved.	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	Reserved.	101b	Reserved.																																		
010b	Divide-by 4	110b	Reserved.																																		
011b	Divide-by 8	111b	Reserved.																																		
27:24	<p>PowerStepUp. Read-write. Cold reset: 0000b. 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 NB steps is clipped to 50 ns. The compute unit steps are not clipped. • BIOS: IF (MultiLink) THEN 0011b ELSE 1000b ENDIF. 	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: See PowerStepUp. 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.																																				
19:18	Reserved.																																				

17:16	<p>LnkPllLock. Read-write. Cold reset: 00b. BIOS: 01b. Specifies the link PLL lock time applied when the link frequency is programmed to change during a link disconnect-reconnect sequence. The reconnect sequence is delayed to ensure that the PLL is locked.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1 us</td> </tr> <tr> <td>01b</td> <td>10 us</td> </tr> <tr> <td>10b</td> <td>100 us</td> </tr> <tr> <td>11b</td> <td>1000 us</td> </tr> </tbody> </table>	Bits	Description	00b	1 us	01b	10 us	10b	100 us	11b	1000 us
Bits	Description										
00b	1 us										
01b	10 us										
10b	100 us										
11b	1000 us										
15	<p>StutterScrubEn: stutter mode scrub enable. Read-write. Cold reset: 0. BIOS: IF (D18F3x58[DramScrub]! =0) THEN 1 ELSE 0 ENDIF. 1=Enable DRAM scrubbing when LDTSTOP_L is deasserted during stutter mode. One scrub request is sent for each LDTSTOP_L deassertion. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E]. Enabled if DRAM scrubbing is enabled; see 2.9.5 [Memory Scrubbers].</p>										
14	<p>CacheFlushImmOnAllHalt: cache flush immediate on all Halt. Read-write. Cold reset: 0. BIOS: 0. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E]. 1=Flush the caches immediately when all cores in a package are in a non-C0 state. Cache flushing must be enabled in all non-C0 states in order for the caches to be flushed.</p>										
13	<p>MTC1eEn: message triggered C1E enable. Read-write. Cold reset: 0. BIOS: 1. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E]. 1=Enables message triggered C1E.</p>										
12	<p>ClkRampHystCtl: clock ramp hysteresis control. Read-write. Cold reset: 0. BIOS: 0. Specifies the time base for ClkRampHystSel when (D18F4x128[CoreCstateMode] ? (D18F3x[84:80][CpuPrbEn]==0) : (D18F4x11[C:8][CpuPrbEn]==0)). 0=320 ns. 1=1.28 us.</p>										
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 (D18F4x128[CoreCstateMode] ? (D18F3x[84:80][CpuPrbEn]==0) : (D18F4x11[C:8][CpuPrbEn]==0)) 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 (D18F4x128[CoreCstateMode] ? (D18F3x[84:80][CpuPrbEn]==1) : (D18F4x11[C:8][CpuPrbEn]==1)) 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	<p>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 2.5.2.1.2.1 [Software P-state Numbering] .</p>										

D18F3xD8 Clock Power/Timing Control 1

See 2.5.1.5 [Voltage Transitions].

Bits	Description																				
31:28	Reserved.																				
27:24	<p>ReConDel: link reconnect delay. Read-write. Cold reset: 0. BIOS: 3h. Specifies the approximate delay, in us, from the deassertion of LDTSTOP_L until the link initialization process is allowed to start in Gen1 mode if D18F0x[E4,C4,A4,84][LdtStopTriEn]=1 and D18F0x[18C:170][LS2En]=1. The assertion of CTL is delayed until the specified time has elapsed. See 2.12.6 [Link LDTSTOP_L Disconnect-Reconnect] on when this is applied. The receiver is always enabled 1 us after deassertion of LDTSTOP_L, regardless of the setting of this field or other delays in assertion of CTL.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Delay</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>1.6 us</td> </tr> <tr> <td>9h-1h</td> <td><ReConDel+1> us</td> </tr> <tr> <td>Fh-Ah</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Delay	0h	1.6 us	9h-1h	<ReConDel+1> us	Fh-Ah	Reserved												
Bits	Delay																				
0h	1.6 us																				
9h-1h	<ReConDel+1> us																				
Fh-Ah	Reserved																				
23:7	Reserved.																				
6:4	<p>VSRampSlamTime. Read-write. Cold reset: 000b. BIOS: 001b. Specifies the time the processor waits for voltage increases to complete before beginning an additional voltage change or a frequency change.</p> <p>Wait time = (VSRampSlamTime / 12.5mV) * ABS(destination voltage - current voltage).</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>6.25 us</td> <td>100b</td> <td>2.50 us</td> </tr> <tr> <td>001b</td> <td>5.00 us</td> <td>101b</td> <td>1.67 us</td> </tr> <tr> <td>010b</td> <td>4.17 us</td> <td>110b</td> <td>1.25 us</td> </tr> <tr> <td>011b</td> <td>3.13 us</td> <td>111b</td> <td>1.00 us</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	000b	6.25 us	100b	2.50 us	001b	5.00 us	101b	1.67 us	010b	4.17 us	110b	1.25 us	011b	3.13 us	111b	1.00 us
Bits	Description	Bits	Description																		
000b	6.25 us	100b	2.50 us																		
001b	5.00 us	101b	1.67 us																		
010b	4.17 us	110b	1.25 us																		
011b	3.13 us	111b	1.00 us																		
3:0	Reserved.																				

D18F3xDC Clock Power/Timing Control 2

Bits	Description						
31:27	Reserved.						
26	<p>IgnCpuPrbEn: ignore CPU probe enable. Read-write. Cold reset: 0. BIOS: 1. See D18F3x[84:80][CpuPrbEn] and D18F4x11[C:8][CpuPrbEn].</p>						
25:19	<p>CacheFlushOnHaltTmr: cache flush on Halt timer. Read-write. Cold reset: 00h. BIOS: 28h. Specifies how long each core needs to stay in a C-state before it flushes its caches. See CacheFlushOnHaltCtl, D18F3x[84:80][CpuPrbEn], D18F4x128[CoreCstateMode], and D18F4x11[C:8][CacheFlushTmrSel].</p> <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	<p>CacheFlushOnHaltCtl: cache flush on Halt control. Read-write. Cold reset: 000b. BIOS: 111b. Enables cache flush on halt when ((D18F4x128[CoreCstateMode]==1) && (D18F3xDC[CacheFlushOnHaltCtl]!=0)). Specifies what core clock divisor is used after the caches have been flushed, regardless of D18F4x128[CoreCstateMode]. See 2.5.3.3.2 [C-state Cache Flush]. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E].</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>IF (D18F4x128[CoreCstateMode]) THEN Disabled. ELSE Divide-by 1. ENDIF.</td> </tr> <tr> <td>001b</td> <td>Divide-by 2</td> </tr> <tr> <td>010b</td> <td>Divide-by 4</td> </tr> <tr> <td>011b</td> <td>Divide-by 8</td> </tr> <tr> <td>100b</td> <td>Divide-by 16</td> </tr> <tr> <td>101b</td> <td>Divide-by 128</td> </tr> <tr> <td>110b</td> <td>Divide-by 512</td> </tr> <tr> <td>111b</td> <td>Turn off clocks</td> </tr> </tbody> </table>	Bits	Description	000b	IF (D18F4x128[CoreCstateMode]) THEN Disabled. ELSE Divide-by 1. ENDIF.	001b	Divide-by 2	010b	Divide-by 4	011b	Divide-by 8	100b	Divide-by 16	101b	Divide-by 128	110b	Divide-by 512	111b	Turn off clocks
Bits	Description																		
000b	IF (D18F4x128[CoreCstateMode]) THEN Disabled. ELSE Divide-by 1. ENDIF.																		
001b	Divide-by 2																		
010b	Divide-by 4																		
011b	Divide-by 8																		
100b	Divide-by 16																		
101b	Divide-by 128																		
110b	Divide-by 512																		
111b	Turn off clocks																		
15	Reserved.																		
14:12	<p>NbsynPtrAdj: NB/core synchronization FIFO pointer adjust. Read-write; Reset-applied. Cold reset: 000b. BIOS: 101b. 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.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>6h-0h</td> <td>Position the read pointer <NbsynPtrAdj> clock cycles closer to the write pointer</td> </tr> <tr> <td>7h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	6h-0h	Position the read pointer <NbsynPtrAdj> clock cycles closer to the write pointer	7h	Reserved												
Bits	Description																		
6h-0h	Position the read pointer <NbsynPtrAdj> clock cycles closer to the write pointer																		
7h	Reserved																		
11	Reserved.																		
10:8	<p>HwPstateMaxVal: P-state maximum value. Read-write. 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 HwPstateMaxVal=1; if MSRC001_0064 has this bit set and the others do not, then HwPstateMaxVal=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 2.5.2.1.2.2 [Hardware P-state Numbering].</p>																		
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 2.11.2.3 [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												
31:30	IntNodeNum. Specifies the internal node number for multi-node processors. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>internal node 0</td> </tr> <tr> <td>01b</td> <td>internal node 1</td> </tr> <tr> <td>11b-10b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	internal node 0	01b	internal node 1	11b-10b	Reserved				
Bits	Description												
00b	internal node 0												
01b	internal node 1												
11b-10b	Reserved												
29	MultiNodeCpu. 1=The processor is a multi-node processor; G34r1. 0=The processor is a single-node processor; AM3r2 or C32r1.												
28:26	Reserved.												
25	L3Capable. Same-for-all. 1=Specifies that an L3 cache is present. See CPUID Fn8000_0006_EDX .												
24	Reserved.												
23:20	UnGangEn: link unganging enabled. 0=Link is forced into the ganged state and may not be placed into the unganging state. 1=Unganging is supported. See 2.12 [Links] . <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Link 0</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[3]</td> <td>Link 3</td> </tr> </tbody> </table>	Bit	Description	[0]	Link 0	[1]	Link 1	[2]	Link 2	[3]	Link 3		
Bit	Description												
[0]	Link 0												
[1]	Link 1												
[2]	Link 2												
[3]	Link 3												
19	x2Apic: x2APIC capability. Value: 0.												
18:16	MpCap: MP capability. Specifies the maximum number of processors supported. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>8 processors for single-node; 4 processors for dual-node.</td> </tr> <tr> <td>100b-001b</td> <td>Reserved</td> </tr> <tr> <td>101b</td> <td>4 processors.</td> </tr> <tr> <td>110b</td> <td>2 processors.</td> </tr> <tr> <td>111b</td> <td>1 processor.</td> </tr> </tbody> </table>	Bits	Description	000b	8 processors for single-node; 4 processors for dual-node.	100b-001b	Reserved	101b	4 processors.	110b	2 processors.	111b	1 processor.
Bits	Description												
000b	8 processors for single-node; 4 processors for dual-node.												
100b-001b	Reserved												
101b	4 processors.												
110b	2 processors.												
111b	1 processor.												
15	Reserved.												
14	MultVidPlane: multiple VID plane capable. Value: 1.												
13:12	Reserved.												
11	LnkRtryCap: link error-retry capable.												
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.

D18F3xFC CPUID Family/Model/Stepping

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: Product-specific.
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.

D18F3x138 DCT0 Bad Symbol Identification

This register is set by software and read by hardware to manage history-based ECC checks during chipkill conditions. The MCT uses this information during fetches to determine if errors are correctable or uncorrectable. See 2.13.2.4 [Software Managed Bad Symbol Identification].

Bits	Description
31	BadDramAllCsEn0: bad DRAM on all chip selects, DCT0. Read-write. Reset: 0. Indicates that the symbol number specified in BadDramSymbol0 applies to all chip selects. This most likely indicates a fault outside the DRAM, since it affects all ranks.
30:12	Reserved.
11	BadDramCsVal0: bad DRAM chip select valid, DCT0. Read-write. Reset: 0. Indicates that BadDramCs0 and BadDramSymbol0 contain valid information.
10:9	Reserved.
8:4	BadDramSymbol0: bad DRAM symbol, DCT0. Read-write. Reset: 0h. Indicates the bad symbol number within the rank. Values 0h through Fh correspond to data byte 0h through Fh, respectively. Value 10h corresponds to ECC byte 0, and value 11h corresponds to ECC byte 1.
3:0	BadDramCs0: bad DRAM chip select, DCT0. Read-write. Reset: 0h. Indicates the chip select value which is known bad. This chip select value identifies the rank in error.

D18F3x13C DCT1 Bad Symbol Identification

See [D18F3x138 \[DCT0 Bad Symbol Identification\]](#).

Bits	Description
31	BadDramAllCsEn1: bad DRAM on all chip selects, DCT1. See: D18F3x138[BadDramAllCsEn0] .
30:12	Reserved.
11	BadDramCsVal1: bad DRAM chip select valid, DCT1. See: D18F3x138[BadDramCsVal0] .
10:9	Reserved.
8:4	BadDramSymbol1: bad DRAM symbol, DCT1. See: D18F3x138[BadDramSymbol0] .
3:0	BadDramCs1: bad DRAM chip select, DCT1. See: D18F3x138[BadDramCs0] .

D18F3x140 SRI to XCS Token Count

Read-write; Reset-applied. [D18F3x140](#), [D18F3x144](#), and [D18F3x1\[54,50,4C,48\]](#) specify the number of XCS (XBAR command scheduler) entries assigned to each virtual channel within each source port. See [2.9 \[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
- Link: 8 * 4 (one group per link).
- Total: 64, which is the total number of entries supported by XCS.

The defaults for [D18F3x140](#), [D18F3x1\[54,50,4C,48\]](#), and [D18F3x158](#) do not allocate any tokens in the isochronous channel. If isochronous flow control mode (IFCM) is enabled ([D18F0x\[E4,C4,A4,84\]\[IsocEn\]](#)) or display refresh mode is enabled ([D18F0x68\[DispRefModeEn\]](#)), then the XCS token counts must be changed.

- If IFCM is enabled on any link, then the [D18F3x140\[IsocReqTok, IsocPreqTok, and IsocRspTok\]](#) must each be non-zero. Or, in display refresh mode, [D18F3x140\[IsocReqTok and IsocRspTok\]](#) must be non-zero. This requires tokens to be reduced elsewhere to avoid exceeding the 64 token maximum. Links which are not connected or links which are ganged include excess tokens which may be used for this purpose.
- If IFCM is enabled on any link, then it may be advantageous to allocate isochronous tokens to that link/sub-link in [D18F3x1\[54,50,4C,48\]](#). However this would result in excessive tokens for a fully populated system, especially if the links are unganged. To account for this, the processor supports IFCM being enabled on a link without allocating dedicated isochronous XCS tokens. In this case:
 - The isochronous channel uses the base channel tokens.
 - The isochronous channel has preferential access to these tokens.
- If an IOMMU is present on a link, [D18F3x1\[54,50,4C,48\]\[IsocReqTok\]](#) for that link must be non-zero.
- In display refresh mode, [D18F3x1\[54,50,4C,48\]\[IsocReqTok\]](#) and [D18F3x1\[54,50,4C,48\]\[IsocPreqTok\]](#) for the enabled link and [D18F3x158\[LnkToXcsDRToken\]](#) must be non-zero.

See [2.9.3.2.5.1 \[Recommended Buffer Count Settings Overview\]](#).

Table 220: BIOS Recommendations for [D18F3x140\[FreeTok, UpRspTok\]](#)

Condition		D18F3x140	
		FreeTok	UpRspTok
SCM		10	3
(MCM1 MCM2h)	PrbFltrEn	9	3
	~PrbFltrEn	10	3

Table 220: BIOS Recommendations for D18F3x140[FreeTok, UpRspTok]

Condition	D18F3x140	
	FreeTok	UpRspTok
MCM2	11	1
MCM4h	10	3
MCM4	9	1

Bits	Description
31:24	Reserved.
23:20	FreeTok: free tokens. Cold reset: Eh. BIOS: Table 220 . 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. Cold reset: 0. BIOS: 1.
15:14	IsocPreqTok: isochronous posted request tokens. Cold reset: 0. BIOS: 0.
13:12	IsocReqTok: isochronous request tokens. Cold reset: 0. BIOS: 1.
11:10	DnRspTok: downstream response tokens. Cold reset: 1. BIOS: 1.
9:8	UpRspTok: upstream response tokens. Cold reset: 3. BIOS: Table 220 .
7:6	DnPreqTok: downstream posted request tokens. Cold reset: 1. BIOS: 1.
5:4	UpPreqTok: upstream posted request tokens. Cold reset: 1. BIOS: 1.
3:2	DnReqTok: downstream request tokens. Cold reset: 1. BIOS: 1.
1:0	UpReqTok: upstream request tokens. Cold reset: 1. BIOS: 1.

D18F3x144 MCT to XCS Token Count

Read-write; Reset-applied. See [D18F3x140](#).

Bits	Description
31:8	Reserved.
7:4	ProbeTok: probe tokens. Cold reset: 2h. BIOS: IF (~PrbFltrEn) THEN 5h ELSE 2h ENDIF.
3:0	RspTok: response tokens. Cold reset: 8h. BIOS: IF (~PrbFltrEn) THEN 5h ELSE 8h ENDIF.

D18F3x1[54,50,4C,48] Link to XCS Token Count

Read-write; Reset-applied.

Table 221: Register Mapping for D18F3x1[54,50,4C,48]

Register	Function
D18F3x148	Link 0
D18F3x14C	Link 1
D18F3x150	Link 2
D18F3x154	Link 3

See [D18F3x140](#). The cold reset default value for some of the fields of this register vary based on the

ganged/unganged state specified by `D18F0x[18C:170][Ganged]`. Most of the fields in this register are duplicated for each sublink; if the link is ganged, then the sublink 0 fields apply and the sublink 1 fields are 0. See 2.12.1.3.2 [Unused Links]. See 2.9.3.2.5.1 [Recommended Buffer Count Settings Overview].

 Table 222: BIOS Recommendations for `D18F3x1[54,50,4C,48]`

Condition		D18F3x1[54,50,4C,48]															
		IsocRspTok1	IsocPreqTok1	IsocReqTok1	ProbeTok1	RspTok1	PReqTok1	ReqTok1	IsocRspTok0	IsocPreqTok0	IsocReqTok0	FreeTok[3:0]	ProbeTok0	RspTok0	PReqTok0	ReqTok0	
~IoLink & LinkGang	SCM	0	0	0	0	0	0	0	0	0	1	3	2	2	2	2	
	(MCM1 MCM2h)	0	0	0	0	0	0	0	0	0	1	0	2	2	2	2	
	(MCM2 MCM4h)	0	0	0	0	0	0	0	0	0	1	0	1	2	2	2	
	MCM4	0	0	0	0	0	0	0	0	0	1	0	2	1	1	2	
~IoLink & ~LinkGang	(MCM1 MCM2h)	~PrbFltrEn	0	0	0	1	1	1	1	0	0	1	0	1	1	1	1
		PrbFltrEn	0	0	0	1	1	1	1	0	0	1	0	1	2	1	1
	MCM2	0	0	1	1	1	1	1	0	0	1	2	1	1	1	1	
	MCM4h	0	0	1	1	1	1	1	0	0	1	4	1	1	1	1	
	MCM4	0	0	1	1	1	1	1	0	0	1	0	1	1	1	1	
IoLink & LinkGang	SCM	0	0	0	0	0	0	0	0	0	1	3	0	2	2	2	
	(MCM1 MCM2h MCM2 MCM4h)	0	0	0	0	0	0	0	0	0	1	0	0	2	2	2	
	MCM4	0	0	0	0	0	0	0	0	0	2	0	2	2	2	2	

Bits	Description
31:30	FreeTok[3:2]: free tokens. See: FreeTok[1:0].
29	Reserved.
28	IsocRspTok1: isochronous response tokens sublink 1. Cold reset: 0. BIOS: Table 222 .
27	Reserved.
26	IsocPreqTok1: isochronous posted request tokens sublink 1. Cold reset: 0. BIOS: Table 222 .
25	Reserved.
24	IsocReqTok1: isochronous request tokens sublink 1. Cold reset: 0. BIOS: Table 222 .
23:22	ProbeTok1: probe tokens sublink 1. Cold reset: IF (LinkGang) THEN 0 ELSE 1 ENDIF. BIOS: Table 222 .
21:20	RspTok1: response tokens sublink 1. Cold reset: IF (LinkGang) THEN 0 ELSE 1 ENDIF. BIOS: Table 222 .

19:18	PReqTok1: posted request tokens sublink 1. Cold reset: IF (LinkGang) THEN 0 ELSE 1 ENDIF. BIOS: Table 222 .
17:16	ReqTok1: request tokens sublink 1. Cold reset: IF (LinkGang) THEN 0 ELSE 1 ENDIF. BIOS: Table 222 .
15:14	FreeTok[1:0]: free tokens. FreeTok[3:0] = {FreeTok[3:2], FreeTok[1:0]}. Cold reset: 0000b. If the link is unganged, the free tokens are shared between the two sublinks. BIOS: Table 222 .
13:12	IsocRspTok0: isochronous response tokens sublink 0. Cold reset: 0. BIOS: Table 222 .
11:10	IsocPreqTok0: isochronous posted request tokens sublink 0. Cold reset: 0. BIOS: Table 222 .
9:8	IsocReqTok0: isochronous request tokens sublink 0. Cold reset: 0. BIOS: Table 222 .
7:6	ProbeTok0: probe tokens sublink 0. Cold reset: IF (LinkGang) THEN 2 ELSE 1 ENDIF. BIOS: Table 222 .
5:4	RspTok0: response tokens sublink 0. Cold reset: IF (LinkGang) THEN 2 ELSE 1 ENDIF. BIOS: Table 222 .
3:2	PReqTok0: posted request tokens sublink 0. Cold reset: IF (LinkGang) THEN 2 ELSE 1 ENDIF. BIOS: Table 222 .
1:0	ReqTok0: request tokens sublink 0. Cold reset: IF (LinkGang) THEN 2 ELSE 1 ENDIF. BIOS: Table 222 .

D18F3x158 Link to XCS Token Count

Reset-applied. See [D18F3x140](#).

Bits	Description
31:4	Reserved.
3:0	LnkToXcsDRToken: display refresh tokens all links. Read-write. Cold reset: 0. BIOS: IF (\sim MultiLink && (UmaDr UmaIpcm)) THEN 3h ELSE 0h ENDIF.

D18F3x160 NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)

See 2.13.1.7 [Error Thresholding]. [D18F3x160](#) is associated with the DRAM error type. See [MSR0000_0413](#).

Bits	Description
31	Valid. Read-only. Reset: 1.
30	CntP: counter present. Read-only. Reset: 1.
29	Locked. Read-only. Reset: 0.
28:24	Reserved.
23:20	LvtOffset: LVT offset. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0h. BIOS: 1h.
19	CntEn: counter enable. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0.
18:17	IntType: interrupt type. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Cold reset: 0.
16	Ovrflw: overflow. IF (Locked) THEN Read-only; set-by-hardware. ELSE Read-write; set-by-hardware. ENDIF. Cold reset: 0.

15:12	Reserved.
11:0	ErrCnt: error counter. IF (Locked) THEN Read-only; updated-by-hardware. ELSE Read-write; updated-by-hardware. ENDIF. Cold reset: 0.

D18F3x168 NB Machine Check Misc (Link Thresholding) 1 (MC4_MISC1)

See 2.13.1.7 [Error Thresholding]. D18F3x168 is associated with the link error type. See MSRC000_0408.

Bits	Description
31	Valid. Read-only. Reset: 1.
30	CntP: counter present. Read-only. Reset: 1.
29	Locked. Read-only. Reset: 0.
28:24	Reserved.
23:20	LvtOffset: LVT offset. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0h. BIOS: 1h.
19	CntEn: counter enable. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0.
18:17	IntType: interrupt type. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Cold reset: 0.
16	Ovrflw: overflow. IF (Locked) THEN Read-only; set-by-hardware. ELSE Read-write; set-by-hardware. ENDIF. Cold reset: 0.
15:12	Reserved.
11:0	ErrCnt: error counter. IF (Locked) THEN Read-only; updated-by-hardware. ELSE Read-write; updated-by-hardware. ENDIF. Cold reset: 0.

D18F3x170 NB Machine Check Misc (L3 Thresholding) 2 (MC4_MISC2)

See 2.13.1.7 [Error Thresholding]. D18F3x170 is associated with the L3 error type. See MSRC000_0409.

Bits	Description
31	Valid. Read-only. Reset: 1.
30	CntP: counter present. Read-only. Reset: 1.
29	Locked. Read-only. Reset: 0.
28:24	Reserved.
23:20	LvtOffset: LVT offset. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0h. BIOS: 1h.
19	CntEn: counter enable. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0.
18:17	IntType: interrupt type. IF (Locked) THEN Read-only. ELSE Read-write. ENDIF. Cold reset: 0.
16	Ovrflw: overflow. IF (Locked) THEN Read-only; set-by-hardware. ELSE Read-write; set-by-hardware. ENDIF. Cold reset: 0.
15:12	Reserved.
11:0	ErrCnt: error counter. IF (Locked) THEN Read-only; updated-by-hardware. ELSE Read-write; updated-by-hardware. ENDIF. Cold reset: 0.

D18F3x17C Extended Freelist Buffer Count

Bits	Description
31:4	Reserved.
3:0	SPQPrbFreeCBC: XBAR to SRI Probe command buffer freelist. Read-write. Cold reset: 9h. Freelist entries are shared between L3 and XBAR.

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	EccSymbolSize: ECC symbol size and code selection. Read-write. BIOS: See 2.13.2 [DRAM Considerations for ECC] . 1=x8 symbol size and code used. 0=x4 symbol size and code used.
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=Enable sync flood 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=Enable sync flood 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=Enable sync flood 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. The state of this field is ignored if SyncFloodOnUsPwDatErr==1.
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. The state of this field is ignored if SyncFloodOnUsPwDatErr==1.
17:11	Reserved.

10	Reserved.
9	SyncFloodOnUCNbAry: sync flood on UC NB array error. Read-write. BIOS: 1. 1=Enable sync flood on detection of an UC error in an NB array.
8	SyncFloodOnProtErr: sync flood on protocol error. Read-write. BIOS: 1. 1=Enable sync flood on detection of link protocol error, L3 protocol error, and probe filter 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	DisPciCfgCpuMstAbtRsp. 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	SyncFloodOnUsPwDatErr: sync flood on upstream posted write data error. Read-write. BIOS: 1. 1=Enable sync flood generation if McaLogUsPwDataErrEn = 1 and an upstream posted write data error is detected; setting of PwP2pDatErrRmtPropDis and PwP2pDatErrLclPropDis are ignored.
0	McaLogUsPwDatErrEn: MCA log of upstream posted write data error enable. Read-write. BIOS:1. 1=Enable logging of upstream posted write data errors in MCA (if NB MCA registers are appropriately enabled and configured).

D18F3x188 NB Configuration 2 (NB_CFG2)

Same-for-all.

Bits	Description
31:10	Reserved.
9	DisL3HiPriFreeListAlloc. Read-write. Reset: 0. BIOS: 1. 1=Disables normal SRQ entry scheme which gives higher priority to L3 than XBAR.
8:0	Reserved.

D18F3x190 Downcore Control

Cold reset: 0000_0000h. See [2.4.3 \[Processor Cores and Downcoring\]](#) and [2.4.3.1 \[Software Downcoring using D18F3x190\[DisCore\]\]](#).

Bits	Description
31:0	DisCore. Read-write; reset-applied. 1=Disable core. 0=Core enabled. [0]=Core 0; ...; [N]=Core N.

D18F3x1A0 Core Interface Buffer Count

- The following buffer allocations rules must be satisfied:
 - CpuCmdBufCnt rules:
 - IF (D18F3xE8[L3Capable]==0) D18F3x1A0[CpuCmdBufCnt] >= 2.
 - IF (D18F3xE8[L3Capable]==1) D18F3x1A0[CpuCmdBufCnt] >= 1.
 - L3FreeListCBC <= (L3qSize - (CpuCmdBufCnt * NumOfCompUnitsOnNode)).
 - L3FreeListCBC >= (CpuToNbFreeBufCnt * NumOfCompUnitsOnNode).
 - L3ToSriReqCBC rules:
 - IF (D18F3xE8[L3Capable]==1) THEN (D18F3x1A0[L3ToSriReqCBC] >= NumOfCompUnitsOnNode).
 - IF (D18F3xE8[L3Capable]==0) D18F3x1A0[L3ToSriReqCBC] >= 0.

Bits	Description																
31:18	Reserved.																
17:16	<p>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.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Reserved.</td> </tr> <tr> <td>01b</td> <td>1 token</td> </tr> <tr> <td>10b</td> <td>2 tokens</td> </tr> <tr> <td>11b</td> <td>3 tokens</td> </tr> </tbody> </table>	Bits	Description	00b	Reserved.	01b	1 token	10b	2 tokens	11b	3 tokens						
Bits	Description																
00b	Reserved.																
01b	1 token																
10b	2 tokens																
11b	3 tokens																
15	Reserved.																
14:12	<p>L3ToSriReqCBC: L3 cache to SRI request command buffer count. Read-write; Reset-applied. Cold reset: 4h. BIOS: 4h.</p>																
11:9	Reserved.																
8:4	<p>L3FreeListCBC: L3 free list command buffer counter for compute unit requests. Read-write; Reset-applied. Cold reset: Product-specific. BIOS: IF (NumOfCompUnitsOnNode==1) THEN 1Ch ELSEIF (NumOfCompUnitsOnNode==2) THEN 18h ELSEIF (NumOfCompUnitsOnNode==3) THEN 14h ELSEIF (NumOfCompUnitsOnNode==4) THEN 10h ELSEIF (NumOfCompUnitsOnNode==5) THEN 11h ELSE 0Eh ENDIF.</p> <table border="1"> <thead> <tr> <th>Compute Units</th> <th>Cold reset</th> <th>Compute Units</th> <th>Cold reset</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1Ch</td> <td>2</td> <td>18h</td> </tr> <tr> <td>3</td> <td>14h</td> <td>4</td> <td>10h</td> </tr> <tr> <td>5</td> <td>Reserved</td> <td>6</td> <td>Reserved</td> </tr> </tbody> </table>	Compute Units	Cold reset	Compute Units	Cold reset	1	1Ch	2	18h	3	14h	4	10h	5	Reserved	6	Reserved
Compute Units	Cold reset	Compute Units	Cold reset														
1	1Ch	2	18h														
3	14h	4	10h														
5	Reserved	6	Reserved														
3	Reserved.																
2:0	<p>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: Product-specific. BIOS: 4.</p>																

D18F3x1B8 L3 Control 1

This register is reserved if the processor does not include an L3 cache. This register controls various performance modes of the L3 cache.

Block Aggressive Neighbors (BAN) is a mechanism of preventing a poor cache behaving application from displacing cache lines of good cache behaving applications in a multi-compute-unit environment. This mode is

enabled by the L3BanMode field. The mode works by computing a figure of merit (FOM) per compute-unit and an average FOM of all compute-units. The FOM is a function of hits per L3 request and L3 requests per cycle. When the ratio between the FOM for a compute-unit and the average FOM of all compute-units falls below a particular level, one of two BAN modes is engaged: full BAN and half BAN. In full BAN mode, L2 victims from that compute unit are allocated at the LRU position instead of MRU. In half BAN mode, L2 victims from that compute unit are allocated at halfway between LRU and MRU. The L3HalfBanLvl and L3FullBanLvl control bits are used to set thresholds for entering full and half BAN modes. For example the default setting for L3FullBanLvl is 0. Once the compute-unit FOM falls below 1/16 of the average FOM, full ban mode is enabled.

Bits	Description										
31:29	Reserved.										
28	ProtoErrHalt: Halt on protocol error. Read-write. Reset: 0. 1=When an L3 protocol error is detected freeze the request queue (Halt request processing); block all L3 MCA error logging/reporting (but not error action). This state, when entered, can only be exited by a warm reset.										
27	L3ATMMModeEn. Read-write; Same-for-all. Reset: 0. BIOS: See 2.9.4.2. 1= Enable Accelerated Transition to Modified protocol in L3. This mode enables usage of a new state MuW (ModifiedUnWritten) to the existing MOESI protocol. Must be programmed to the same state as D18F0x68[ATMMModeEn].										
26:23	Reserved.										
22	L3HalfBanLvl: L3 block aggressive neighbors half ban level. Read-write. Reset: 0. Specifies the FOM ratio to allocate an entry halfway between LRU and MRU position. 0=1/8 FOM ratio. 1=1/32 FOM ratio.										
21	L3FullBanLvl: L3 block aggressive neighbors full ban level. Read-write. Reset: 0. Specifies the FOM ratio to allocate an entry at LRU position. 0=1/32 FOM ratio. 1=Disabled.										
20:19	L3BanMode: L3 block aggressive neighbors mode. Read-write. Reset: 10b. Enables BAN mode and specifies the duty cycle size. A duty cycle is divided into 8 slices with BAN active on 7 slices and disabled on 1 slice. Must be disabled if L3 cache partitioning is enabled; See 2.9.4.3 [L3 Cache Partitioning]. <table border="1" data-bbox="267 1249 730 1417"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Disabled</td> </tr> <tr> <td>01b</td> <td>2 million cycles</td> </tr> <tr> <td>10b</td> <td>4 million cycles</td> </tr> <tr> <td>11b</td> <td>16 million cycles</td> </tr> </tbody> </table>	Bits	Description	00b	Disabled	01b	2 million cycles	10b	4 million cycles	11b	16 million cycles
Bits	Description										
00b	Disabled										
01b	2 million cycles										
10b	4 million cycles										
11b	16 million cycles										
18:13	Reserved.										
12	L3PrivReplEn: L3 private replacement enable. Read-write. Reset: 0. 1=A line that is not shared with other cores, is invalidated from the cache when a read hit occurs to that line from the core that originally allocated the line, regardless of the state of its preference bit. The line is returned to the allocating core in its original state. This bit should be set to 1 by BIOS.										
11:5	Reserved.										
4	L3ScrbRedirDis: L3 scrubber redirect disable. Read-write. Reset: 0. BIOS: 2.10.5.6. 0=A correctable ECC data error causes the scrubber to be redirected to correct the error. This does not require sequential scrubbing to be enabled. 1=A correctable ECC data error is left in place until the sequential scrubber reaches the location of the error(s).										
3:0	Reserved.										

D18F3x1C4 L3 Cache Parameter

D18F3x1C4 is reserved if (D18F3xE8[L3Capable]==0).

Bits	Description																
31	L3TagInit: L3 tag initialization. Read-write; cleared-by-hardware. Reset: 0. 1=Initialize the L3 cache tag arrays. 0=L3 cache tag initialization is complete. This bit is cleared by hardware when the tag initialization is complete. This bit should not be written while initialization is in progress.																
30:16	Reserved.																
15:12	L3SubcacheSize3: L3 subcache size 3. See: L3SubcacheSize0. Reset: Product-specific.																
11:8	L3SubcacheSize2: L3 subcache size 2. See: L3SubcacheSize0. Reset: Product-specific.																
7:4	L3SubcacheSize1: L3 subcache size 1. See: L3SubcacheSize0. Reset: Product-specific.																
3:0	L3SubcacheSize0: L3 subcache size 0. Read-only. Reset: Product-specific. Specifies the size of L3 subcache 0. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Bh-0h</td> <td>Reserved</td> <td>Eh</td> <td>1 MB</td> </tr> <tr> <td>Ch</td> <td>2 MB</td> <td>Fh</td> <td>0 MB</td> </tr> <tr> <td>Dh</td> <td>1 MB</td> <td></td> <td></td> </tr> </tbody> </table>	Bits	Description	Bits	Description	Bh-0h	Reserved	Eh	1 MB	Ch	2 MB	Fh	0 MB	Dh	1 MB		
Bits	Description	Bits	Description														
Bh-0h	Reserved	Eh	1 MB														
Ch	2 MB	Fh	0 MB														
Dh	1 MB																

D18F3x1CC IBS Control

Reset: 0000_0000h. 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						

D18F3x1D4 Probe Filter Control

See 2.9.4.1 [Probe Filter].

Bits	Description
31:30	Reserved.
29	PFLoIndexHashEn: probe filter low index hash enable. Read-write; Same-for-all. Reset: 0. BIOS: See 2.9.4.2. 1=Include address bits[8:12] in the subcache selection hash. 0=Do not include address bits[8:12] in the subcache selection hash. The subcache selection hash is computed using an exclusive OR of system address bits.

28	PFEccError: probe filter ECC error. Set-by-hardware; write-1-to-clear. Reset: 0. 1=An ECC error was encountered on a directory read.										
27:24	LvtOffset: probe filter error interrupt LVT offset. Read-write; Same-for-all. Reset: 0h. This 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]).										
23:22	PFErrInt[1:0]: probe filter error interrupt type. Read-write; Same-for-all. Reset: 00b. <table border="1"> <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 LVT based interrupt</td> </tr> <tr> <td>10b</td> <td>SMI</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	no interrupt	01b	APIC LVT based interrupt	10b	SMI	11b	Reserved
Bits	Description										
00b	no interrupt										
01b	APIC LVT based interrupt										
10b	SMI										
11b	Reserved										
21:20	PFPreferredSOREpl: PF preferred SO replacement mode. Read-write; Same-for-all. Reset: 00b. BIOS: See 2.9.4.2 . Preferentially retain entries in S and O states and replace entries in M/E or S1 states. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Use preferredSO replacement 14/16 times.</td> </tr> <tr> <td>01b</td> <td>Use preferredSO replacement 15/16 times.</td> </tr> <tr> <td>10b</td> <td>Use preferredSO replacement 12/16 times.</td> </tr> <tr> <td>11b</td> <td>Disable preferredSO replacement.</td> </tr> </tbody> </table>	Bits	Description	00b	Use preferredSO replacement 14/16 times.	01b	Use preferredSO replacement 15/16 times.	10b	Use preferredSO replacement 12/16 times.	11b	Disable preferredSO replacement.
Bits	Description										
00b	Use preferredSO replacement 14/16 times.										
01b	Use preferredSO replacement 15/16 times.										
10b	Use preferredSO replacement 12/16 times.										
11b	Disable preferredSO replacement.										
19	PFInitDone: probe filter initialization complete. Read-only; updated-by-hardware. Reset: 0. 1=Probe filter initialization is complete. 0=Probe filter is not enabled or initialization is in progress.										
18	Reserved.										
17	PFWayHashEn. Read-write; Same-for-all. Reset: 0. BIOS: See 2.9.4.2 . 1=Select hash as described by the following table. The hash that gets applied is a function of PFWayNum. <table border="1"> <thead> <tr> <th>PFWayNum</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1-way (SubCache way [15]) allocated to probe filter.</td> </tr> <tr> <td>01b</td> <td>2-ways (SubCache ways [15:14]) allocated to probe filter, as selected by Way[0]. IF (PFWayHashEn==0) THEN Way[0]=PhysAddr[20]. ELSE Way[0]=(PhysAddr[20] ^ PhysAddr[27] ^ PhysAddr[29] ^ PhysAddr[32]). ENDIF.</td> </tr> <tr> <td>10b</td> <td>4-ways (SubCache ways [15:12]) allocated to probe filter, as selected by Way[1:0]. IF (PFWayHashEn==0) THEN Way[1:0]=PhysAddr[21:20]. ELSE Way[1:0]=(PhysAddr[21] ^ PhysAddr[28] ^ PhysAddr[33]); Way[0]=(PhysAddr[20] ^ PhysAddr[27] ^ PhysAddr[29] ^ PhysAddr[32]). ENDIF.</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	PFWayNum	Description	00b	1-way (SubCache way [15]) allocated to probe filter.	01b	2-ways (SubCache ways [15:14]) allocated to probe filter, as selected by Way[0]. IF (PFWayHashEn==0) THEN Way[0]=PhysAddr[20]. ELSE Way[0]=(PhysAddr[20] ^ PhysAddr[27] ^ PhysAddr[29] ^ PhysAddr[32]). ENDIF.	10b	4-ways (SubCache ways [15:12]) allocated to probe filter, as selected by Way[1:0]. IF (PFWayHashEn==0) THEN Way[1:0]=PhysAddr[21:20]. ELSE Way[1:0]=(PhysAddr[21] ^ PhysAddr[28] ^ PhysAddr[33]); Way[0]=(PhysAddr[20] ^ PhysAddr[27] ^ PhysAddr[29] ^ PhysAddr[32]). ENDIF.	11b	Reserved
PFWayNum	Description										
00b	1-way (SubCache way [15]) allocated to probe filter.										
01b	2-ways (SubCache ways [15:14]) allocated to probe filter, as selected by Way[0]. IF (PFWayHashEn==0) THEN Way[0]=PhysAddr[20]. ELSE Way[0]=(PhysAddr[20] ^ PhysAddr[27] ^ PhysAddr[29] ^ PhysAddr[32]). ENDIF.										
10b	4-ways (SubCache ways [15:12]) allocated to probe filter, as selected by Way[1:0]. IF (PFWayHashEn==0) THEN Way[1:0]=PhysAddr[21:20]. ELSE Way[1:0]=(PhysAddr[21] ^ PhysAddr[28] ^ PhysAddr[33]); Way[0]=(PhysAddr[20] ^ PhysAddr[27] ^ PhysAddr[29] ^ PhysAddr[32]). ENDIF.										
11b	Reserved										
16	DisDirectedPrb: disable directed probes. Read-write. Reset: 0. 1=Force broadcast probes on directory hits.										
15:12	PFSubCacheEn: probe filter subcache enable. Read-write; Same-for-all. Reset: 0h. BIOS: See 2.9.4.2 . Each bit applies to one of the four L3 subcaches. 1=Enable allocation of the probe filter directory within the subcache. 1, 2 or 4 subcaches can be enabled. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Subcache 0 enable</td> </tr> <tr> <td>[1]</td> <td>Subcache 1 enable</td> </tr> <tr> <td>[2]</td> <td>Subcache 2 enable</td> </tr> <tr> <td>[3]</td> <td>Subcache 3 enable</td> </tr> </tbody> </table>	Bit	Description	[0]	Subcache 0 enable	[1]	Subcache 1 enable	[2]	Subcache 2 enable	[3]	Subcache 3 enable
Bit	Description										
[0]	Subcache 0 enable										
[1]	Subcache 1 enable										
[2]	Subcache 2 enable										
[3]	Subcache 3 enable										
11:10	PFSubCacheSize3: probe filter subcache 3 size. See: PFSubCacheSize2.										

9:8	PFSubCacheSize2: probe filter subcache 2 size. Read-write; Same-for-all. Reset: 00b. BIOS: See 2.9.4.2. This specifies the size of the selected L3 subcache used by the probe filter. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1 MB</td> </tr> <tr> <td>01b</td> <td>2 MB</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 MB	01b	2 MB	10b	Reserved	11b	Reserved
<u>Bits</u>	<u>Description</u>										
00b	1 MB										
01b	2 MB										
10b	Reserved										
11b	Reserved										
7:6	PFSubCacheSize1: probe filter subcache 1 size. See: PFSubCacheSize0.										
5:4	PFSubCacheSize0: probe filter subcache 0 size. Read-write; Same-for-all. Reset: 00b. BIOS: See 2.9.4.2. The size of the selected L3 subcache used by the probe filter. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1 MB</td> </tr> <tr> <td>01b</td> <td>2 MB</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 MB	01b	2 MB	10b	Reserved	11b	Reserved
<u>Bits</u>	<u>Description</u>										
00b	1 MB										
01b	2 MB										
10b	Reserved										
11b	Reserved										
3:2	PFWayNum: probe filter way number. Read-write; Same-for-all. Reset: 00b. BIOS: See 2.9.4.2. This specifies the number of ways in selected L3 subcache(s) allocated for the probe filter directory. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1-way</td> </tr> <tr> <td>01b</td> <td>2-way</td> </tr> <tr> <td>10b</td> <td>4-way</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	1-way	01b	2-way	10b	4-way	11b	Reserved
<u>Bits</u>	<u>Description</u>										
00b	1-way										
01b	2-way										
10b	4-way										
11b	Reserved										
1:0	PFMode: probe filter mode. Read-write; Same-for-all. Reset: 00b. BIOS: See 2.9.4.2. ENDIF. See 2.9.4.1 [Probe Filter] and PrbFltrEn. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Disabled.</td> </tr> <tr> <td>01b</td> <td>Reserved.</td> </tr> <tr> <td>10b</td> <td>EPF4. Probe filter enabled, 4-way.</td> </tr> <tr> <td>11b</td> <td>EPF8. Probe filter enabled, 8-way.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Disabled.	01b	Reserved.	10b	EPF4. Probe filter enabled, 4-way.	11b	EPF8. Probe filter enabled, 8-way.
<u>Bits</u>	<u>Description</u>										
00b	Disabled.										
01b	Reserved.										
10b	EPF4. Probe filter enabled, 4-way.										
11b	EPF8. Probe filter enabled, 8-way.										

D18F3x1E4 SBI Control

This register specifies the behavior associated with the SIC and SID pins which may be used to support SMBus-based sideband interface (SBI) protocol. See 2.13.4 [Sideband Interface (SBI)]. D18F3x1E4, with the exception of SbiAddr, should only be programmed on internal node 0 (D18F3xE8[IntNodeNum]=00b) if D18F3xE8[MultiNodeCpu]=1.

Bits	Description
31	SbiRegWrDn: SBI register write complete. Read-only; updated-by-hardware. Reset: 1. 1=Write to the SBI registers through D18F3x1EC has completed. 0=Write to the SBI registers in progress.
30:12	Reserved.
11:8	LvtOffset: local vector table offset. Read-write. Cold reset: 0000b. BIOS: 3h. This 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]).
7	Reserved.

6:4	SbiAddr: SMBus-based sideband interface address. Read-write. Specifies bits[3:1] of the SMBus address of the processor SBI ports. SMBus address bits [3:1] = {~SA[2],SA[1:0]}. SbiAddr must be unique for all processors on the same SMBus segment. Cold reset: For internal node 0, specified by the SA[2:0] strap pins (value matches the pins until the deassertion of RESET_L for a cold reset only; value is not changed by a warm reset); 000b for internal node 1 or products that do not include SA[2:0] pins. BIOS: It is recommended that BIOS program SbiAddr on internal node 1 to the same value as internal node 0 of the processor.
3	Reserved.
2	Reserved.
1	SbRmiDis: SMBus-based sideband remote management interface disable. Read-only. Reset: Product-specific. 1=The processor does not support SMBus-based SB-RMI protocol.
0	Reserved.

D18F3x1E8 SBI Address

Cold reset: 0000_0000h. See 1.2 [Reference Documents] for the SB-TSI and APML specs.

The SB-RMI (APML) registers can be directly accessed by the processor using D18F3x1E8 and D18F3x1EC. Access to these registers is accomplished as follows:

- Reads:
 - Write D18F3x1E8[SbiBankSel, SbiRegAddr].
 - Read D18F3x1EC.
- Writes:
 - Write D18F3x1E8[SbiBankSel, SbiRegAddr].
 - Write D18F3x1EC.

D18F3x1E8 and D18F3x1EC should only be programmed on internal node 0 (D18F3xE8[IntNodeNum]=00b) if D18F3xE8[MultiNodeCpu]=1. Accesses to reserved offsets result in undefined behavior.

Bits	Description										
31:11	Reserved.										
10:9	SbiByteCnt: SBI byte count. Read-write. Specifies the number consecutive SBI registers to read or write using D18F3x1EC. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1 register</td> </tr> <tr> <td>01b</td> <td>2 registers</td> </tr> <tr> <td>10b</td> <td>3 registers</td> </tr> <tr> <td>11b</td> <td>4 registers</td> </tr> </tbody> </table>	Bits	Description	00b	1 register	01b	2 registers	10b	3 registers	11b	4 registers
Bits	Description										
00b	1 register										
01b	2 registers										
10b	3 registers										
11b	4 registers										
8	SbiBankSel: SBI register bank select. Read-write. Specifies if the SB-RMI registers are accessed. 0=Reserved. 1=SB-RMI registers.										
7:0	SbiRegAddr: SBI SMBus register address. Read-write. Specifies the 8-bit address of the SB-RMI register to access.										

D18F3x1EC SBI Data

Cold reset: 0000_0000h. See D18F3x1E8.

Bits	Description
31:24	SbiRegDat3: SBI SMBus register 3 data. Read-write. Specifies the data to be read or written to the SBI register selected by D18F3x1E8[SbiRegAddr+3] .
23:16	SbiRegDat2: SBI SMBus register 2 data. Read-write. Specifies the data to be read or written to the SBI register selected by D18F3x1E8[SbiRegAddr+2] .
15:8	SbiRegDat1: SBI SMBus register 1 data. Read-write. Specifies the data to be read or written to the SBI register selected by D18F3x1E8[SbiRegAddr+1] .
7:0	SbiRegDat0: SBI SMBus register 0 data. Read-write. Specifies the data to be read or written to the SBI register selected by D18F3x1E8[SbiRegAddr] .

D18F3x1EC_x100 SB-RMI Revision

Bits	Description
7:0	Revision. Read-only. Value: 03h. Specifies the APML revision to which this processor is compliant.

D18F3x1FC Product Information Register 1

Bits	Description
31	Reserved.
30:24	Reserved.
23	Reserved.
22	IF (PROC >= OR_C0) THEN DiDtCfg5. Value: Product-specific. See MSRC001_1028[DiDtCfg5] . ELSE Reserved. ENDIF.
22	Reserved.
21	Reserved.
20	Reserved.
19:17	IF (PROC >= OR_C0) THEN DiDtCfg4. Value: Product-specific. See MSRC001_1028[DiDtCfg4] . ELSE Reserved. ENDIF.
16	IF (PROC < OR_C0) THEN DiDtCfg3. Value: Product-specific. See MSRC001_1028[DiDtCfg3] . ELSE Reserved. ENDIF.
19:16	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] .

3.7 Device [1F:18]h Function 4 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.8 [Configuration Space].

D18F4x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID . Read-only. Value: 1604h.
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: Product-specific. Specifies the offset of the link capabilities block based on which links are supported and unganged.																						
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>All supported links are ganged.</td> </tr> <tr> <td>7Fh-01h</td> <td>Reserved</td> </tr> <tr> <td>80h</td> <td>If link 0 is supported and unganged.</td> </tr> <tr> <td>9Fh-81h</td> <td>Reserved</td> </tr> <tr> <td>A0h</td> <td>If link 0 is ganged/unsupported and link 1 is supported and unganged.</td> </tr> <tr> <td>BFh-A1h</td> <td>Reserved</td> </tr> <tr> <td>C0h</td> <td>If links 0 and 1 are ganged/unsupported and link 2 is supported and unganged.</td> </tr> <tr> <td>DFh-C1h</td> <td>Reserved</td> </tr> <tr> <td>E0h</td> <td>If links 0, 1, and 2 are ganged/unsupported and link 3 is supported and unganged.</td> </tr> <tr> <td>FFh-E1h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00h	All supported links are ganged.	7Fh-01h	Reserved	80h	If link 0 is supported and unganged.	9Fh-81h	Reserved	A0h	If link 0 is ganged/unsupported and link 1 is supported and unganged.	BFh-A1h	Reserved	C0h	If links 0 and 1 are ganged/unsupported and link 2 is supported and unganged.	DFh-C1h	Reserved	E0h	If links 0, 1, and 2 are ganged/unsupported and link 3 is supported and unganged.	FFh-E1h	Reserved
Bits	Description																						
00h	All supported links are ganged.																						
7Fh-01h	Reserved																						
80h	If link 0 is supported and unganged.																						
9Fh-81h	Reserved																						
A0h	If link 0 is ganged/unsupported and link 1 is supported and unganged.																						
BFh-A1h	Reserved																						
C0h	If links 0 and 1 are ganged/unsupported and link 2 is supported and unganged.																						
DFh-C1h	Reserved																						
E0h	If links 0, 1, and 2 are ganged/unsupported and link 3 is supported and unganged.																						
FFh-E1h	Reserved																						

D18F4x[E0,C0,A0,80] Sublink 1 Capability

See [2.12.1.3.1 \[Link Specific Registers\]](#).

Bits	Description
31:0	See: D18F0x[E0,C0,A0,80] .

D18F4x[E4,C4,A4,84] Sublink 1 Control

See [2.12.1.3.1 \[Link Specific Registers\]](#).

Bits	Description
31:0	See: D18F0x[E4,C4,A4,84] .

D18F4x[E8,C8,A8,88] Sublink 1 Frequency/Revision

See [2.12.1.3.1 \[Link Specific Registers\]](#).

Bits	Description
31:0	See: D18F0x[E8,C8,A8,88] .

D18F4x[EC,CC,AC,8C] Sublink 1 Feature Capability

See [2.12.1.3.1 \[Link Specific Registers\]](#).

Bits	Description
31:0	See: D18F0x[EC,CC,AC,8C] .

D18F4x[F0,D0,B0,90] Sublink 1 Base Channel Buffer Count

See 2.12.1.3.1 [Link Specific Registers].

Bits	Description
31:0	See: D18F0x[F0,D0,B0,90].

D18F4x[F4,D4,B4,94] Sublink 1 Isochronous Channel Buffer Count

See 2.12.1.3.1 [Link Specific Registers].

Bits	Description
31:0	See: D18F0x[F4,D4,B4,94].

D18F4x[F8,D8,B8,98] Sublink 1 Link Type

See 2.12.1.3.1 [Link Specific Registers].

Bits	Description
31:0	See: D18F0x[F8,D8,B8,98].

D18F4x[FC,DC,BC,9C] Sublink 1 Link Frequency Extension

See 2.12.1.3.1 [Link Specific Registers].

Bits	Description
31:0	See: D18F0x[FC,DC,BC,9C].

D18F4x104 TDP Accumulator Divisor Control

Bits	Description										
31:14	Reserved.										
13:2	TdpAccDivRate: TDP accumulator divisor rate. Read-write. Reset: 0. BIOS: C8h. Specifies the rate at which the node and compute unit TDP accumulators are effected by TdpAccDivVal. A value of 0 in this field disables the actions specified by TdpAccDivVal.										
1:0	TdpAccDivVal: TDP accumulator divisor value. Read-write. Reset: 0. BIOS: 01b. See TdpAccDivRate. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Divide by 1</td> </tr> <tr> <td>01b</td> <td>Divide by 2</td> </tr> <tr> <td>10b</td> <td>Divide by 4</td> </tr> <tr> <td>11b</td> <td>Reset to 0</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Divide by 1	01b	Divide by 2	10b	Divide by 4	11b	Reset to 0
<u>Bits</u>	<u>Description</u>										
00b	Divide by 1										
01b	Divide by 2										
10b	Divide by 4										
11b	Reset to 0										

D18F4x108 TDP Limit 1

Bits	Description
31:0	Reserved.

D18F4x10C TDP Limit 2

Bits	Description
31:12	Reserved.
11:0	NodeTdpLimit. Read-write; Same-for-all. 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 2.5.2.1.1.1 [TDP Limiting] .

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. Reset: Product-specific. Specifies the minimum amount of time required between TDP-initiated P-state transitions. The minimum amount of time is defined as $\text{MinResTmr} * (\text{CSampleTimer} + 1)$.
12	Reserved.
11:0	CSampleTimer. Read-write. Reset: 0. BIOS: 001h. The ApmSampleTimer and FreeRunSampleTimer rate is $5.12 \text{ us} * (\text{D18F4x110}[\text{CSampleTimer}] + 1)$. The ApmSampleTimer and FreeRunSampleTimer count down from CSampleTimer to 0 before resetting back to CSampleTimer. See 2.5.2.1.1 [Application Power Management (APM)] . The ApmSampleTimer is enabled if $((\text{D18F4x110}[\text{CSampleTimer}] \neq 0) \&\& (\text{D18F4x15C}[\text{ApmMasterEn}] = 1))$. The FreeRunSampleTimer is enabled if $(\text{D18F4x110}[\text{CSampleTimer}] \neq 0)$.

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 \[C-states\]](#).

- [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.

Bits	Description
31:25	Reserved.
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. See. 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E] .
16	CpuPrbEnCstAct1 . See: CpuPrbEnCstAct0. BIOS: 1.
15:9	Reserved.
8	PwrGateEnCstAct0: power gate enable . BIOS: 0. 1=Core power gating is enabled. CacheFlushEnCstAct0 is required to be set if this bit is set. See 2.5.3.3.3 [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"> <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>Divide-by 128</td> </tr> <tr> <td>010b</td> <td>Divide-by 4</td> <td>110b</td> <td>Divide-by 512</td> </tr> <tr> <td>011b</td> <td>Divide-by 8</td> <td>111b</td> <td>Turn off clocks.</td> </tr> </tbody> </table> <p>See CacheFlushTmrSelCstAct0 and 2.9.5 [Memory Scrubbers].</p>	Bits	Description	Bits	Description	000b	Divide-by 1	100b	Divide-by 16	001b	Divide-by 2	101b	Divide-by 128	010b	Divide-by 4	110b	Divide-by 512	011b	Divide-by 8	111b	Turn off clocks.
Bits	Description	Bits	Description																		
000b	Divide-by 1	100b	Divide-by 16																		
001b	Divide-by 2	101b	Divide-by 128																		
010b	Divide-by 4	110b	Divide-by 512																		
011b	Divide-by 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). See 2.5.3.3.2 [C-state Cache Flush]. See CpuPrbEnCstAct0.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Cache flush timer</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No timer. (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>	Bits	Cache flush timer	00b	No timer. (0 us)	01b	D18F3xDC[CacheFlushOnHaltTmr]	10b	D18F4x128[CacheFlushTmr]	11b	Reserved										
Bits	Cache flush timer																				
00b	No timer. (0 us)																				
01b	D18F3xDC[CacheFlushOnHaltTmr]																				
10b	D18F4x128[CacheFlushTmr]																				
11b	Reserved																				
1	<p>CacheFlushEnCstAct0: cache flush enable. BIOS: 1. 1=Enable cache flush if (D18F4x128[CoreCstateMode]==0). See 2.5.3.3.2 [C-state Cache Flush]. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E].</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 and ClkDivisorCstAct0!=000b, 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 or ClkDivisorCstAct0=000b, 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
31:9	Reserved.
8	PwrGateEnCstAct2 . See: D18F4x11[C:8][PwrGateEnCstAct0] . BIOS: 0.
7:5	ClkDivisorCstAct2 . See: D18F4x11[C:8][ClkDivisorCstAct0] . BIOS: 0.
4	Reserved.
3:2	CacheFlushTmrSelCstAct2 . See: D18F4x11[C:8][CacheFlushTmrSelCstAct0] . BIOS: 0.
1	CacheFlushEnCstAct2 . See: D18F4x11[C:8][CacheFlushEnCstAct0] . BIOS: 0.
0	CpuPrbEnCstAct2 . See: D18F4x11[C:8][CpuPrbEnCstAct0] . BIOS: 0.

D18F4x128 C-state Policy Control 1

Reset: 0000_0000h.

Bits	Description						
31	CstateMsgDis: C-state messaging disable . Read-write. Specifies whether any HT Halt entry messages are sent when a core enters a C-state. 0=Halt entry messages are sent. If D18F3xD4[MTC1eEn] = 1, only one Halt entry message is sent from the BSP to the IO hub when all cores in the system are in a non-C0 C-state. 1=Halt entry messages are not sent. See 2.5.3.2 [C-state Request Interface] and 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E] .						
30:21	Reserved.						
20:18	CacheFlushSucMonThreshold: cache flush success monitor threshold . Read-write. BIOS: 000b. 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	CoreStateSaveDestNode: core state save destination node . IF (D18F2x118[LockDramCfg]) THEN Read-only. ELSE Read-write. ENDIF. Specifies the node ID used to save the core state. See 2.10.8 [DRAM CC6 Storage] .						
11:5	CacheFlushTmr: cache flush timer . Read-write. BIOS: 28h. Specifies how long each core needs to stay in a C-state before it flushes its caches. See CoreCstateMode and D18F4x11[C:8][CacheFlushTmrSel] . <table border="1" style="margin-left: 20px;"> <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>(<CacheFlushTmr> * 10.24us) - 5.12us <= Time <= <CacheFlushTmr> * 10.24 us</td> </tr> </tbody> </table>	Bits	Description	00h	<= 5.12 us	7Fh-01h	(<CacheFlushTmr> * 10.24us) - 5.12us <= Time <= <CacheFlushTmr> * 10.24 us
Bits	Description						
00h	<= 5.12 us						
7Fh-01h	(<CacheFlushTmr> * 10.24us) - 5.12us <= Time <= <CacheFlushTmr> * 10.24 us						
4:2	HaltCstateIndex . Read-write. BIOS: 0. Specifies the per-core IO-based C-state that is invoked by a HLT instruction. See CoreCstateMode.						

1	CoreCstatePolicy. Read-write; Same-for-all. 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. See also 2.5.4.1 [Dependencies Between Cores in a Compute Unit] .
0	CoreCstateMode. Read-write; Same-for-all. Specifies whether C-state actions are specified by D18F4x11[C:8] or D18F3x[84:80] and D18F3xDC[CacheFlushOnHaltCtl, CacheFlushOnHaltTmr] . 0=C-state actions are specified by D18F4x11[C:8] . HLT instructions are interpreted as the per-core IO-based C-state pointed to by HaltCstateIndex. 1=C-state actions are specified by D18F3x[84:80] and D18F3xDC[CacheFlushOnHaltCtl, CacheFlushOnHaltTmr] . Per-core IO-based C-state requests are interpreted as Halt special bus cycles and the power management commands specified by F3x84[31:24] and D18F3xDC[CacheFlushOnHaltCtl, CacheFlushOnHaltTmr] are invoked.

D18F4x15C Core Performance Boost Control

Bits	Description
31	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. <ul style="list-style-type: none"> • MSRC001_00[6B:64][CpuFid, CpuDid, CpuVid] • D18F4x10C[NodeTdpLimit] • D18F4x110[MinResTmr] • D18F4x15C[NumBoostStates] • D18F4x16C[CstateCnt, CstateBoost]
30:28	TdpLimitPstate. Read-write. Reset: 0. Specifies the highest performance P-state that has a power consumption less than or equal to the TDP limit. This field is programmed by BIOS and uses software P-state numbering. See 2.5.2.1.2.1 [Software P-state Numbering] and 2.5.2.1.1.2 [Notification of TDP Limit Changes] .
27:8	Reserved.
7	ApmMasterEn: APM master enable. Read-write. Reset: 0. BIOS: 1. 1=Enables the ability to turn on features associated with APM when used in conjunction with the individual feature enable bits. See 2.5.2.1.1 [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 2.5.2.1.1 [Application Power Management (APM)]. See MSRC001_0072[NumBoostStates].										
1:0	BoostSrc: boost source. Read-write. Reset: 0. BIOS: 01b. Specifies whether CPB is enabled or disabled. Under certain conditions boosting is disabled; See 2.5.2.1.9 [Processor-Systemboard Power Delivery Compatibility Check]. <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 processor revision guide for the definition of this field; see 1.2 [Reference Documents].

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 2.5.2.1.2.2 [Hardware P-state Numbering].
5	ApmTdpLimitSts: APM TDP limit status. Read; set-by-hardware; write-1-to-clear. Reset: 0. This bit is set by hardware when D18F5xE8[ApmTdpLimit] changes.
4	ApmTdpLimitIntEn: APM TDP limit interrupt enable. Read-write. Reset: 0. BIOS: 1. 1=Enables the generation of an interrupt using APIC330 of each core for when D18F5xE8[ApmTdpLimit] changes. See 2.5.2.1.1.2 [Notification of TDP Limit Changes].
3:0	Reserved.

D18F4x1[98,90,88,80] Link Phy Offset

Cold reset: 8000_0000h. The links each include an array of registers called F4x1[9C, 94, 8C, 84]_x[NN:00], which are defined following D18F4x1[9C,94,8C,84]. These are used primarily to control link electrical parameters and to program the link BIST engine. D18F4x1[98,90,88,80] [Link Phy Offset] and

D18F4x1[9C,94,8C,84] [Link Phy Data Port] are used to access these registers.

The register number, the number that follows “_x” in the register mnemonic, expands to 16 bit wide and is specified by D18F4x1[98,90,88,80][{UpperLinkPhyOffset, LinkPhyOffset}]. For example, to access D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1] [Link Phy DLL Control], D18F4x1[98,90,88,80][{UpperLinkPhyOffset, LinkPhyOffset}] must be programmed as 530Ah or 520Ah.

The links also include an array of direct map registers, as indicated by the mnemonic suffix “_dm”; alink phy register is not a direct map register if the mnemonic suffix does not end with “_dm”. Direct map registers are accessed by D18F4x1[98,90,88,80][DirectMapEn]=1; non-direct map registers are accessed by D18F4x1[98,90,88,80][DirectMapEn]=0.

Access to these registers is accomplished as follows:

- Reads:
 - Write D18F4x1[98,90,88,80][UpperLinkPhyOffset, LinkPhyOffset, DirectMapEn] and D18F4x1[98,90,88,80][LinkPhyWrite]=0.
 - Poll D18F4x1[98,90,88,80][LinkPhyDone] until it is high.
 - Read the register contents from D18F4x1[9C,94,8C,84].
- Writes:
 - Write all 32 bits of register data to D18F4x1[9C,94,8C,84] (individual byte writes are not supported).
 - Write the register number to D18F4x1[98,90,88,80][UpperLinkPhyOffset, LinkPhyOffset, DirectMapEn] and D18F4x1[98,90,88,80][LinkPhyWrite]=1.
 - Poll D18F4x1[98,90,88,80][LinkPhyDone] until it is high to ensure that the contents of the write have been delivered to the phy.

Read or write accesses to undocumented or undefined register numbers can result in undefined behavior.

Read or write accesses to links that are not implemented on the package complete with undefined results. See 2.12.1.3.2 [Unused Links]

F4x180 and F4x184 are for link 0; F4x188 and F4x18C are for link 1; F4x190 and F4x194 are for link 2; F4x198 and F4x19C are for link 3.

Bits	Description
31	LinkPhyDone: link phy access complete. Read-only. 1=The access to one of the F4x1[9C, 94, 8C, 84]_x[NN:00] registers is complete. 0=The access is still in progress.
30	LinkPhyWrite: link phy read/write select. Read-write. 0=Read one of the F4x1[9C, 94, 8C, 84]_x[NN:00] registers. 1=Write one of the F4x1[9C, 94, 8C, 84]_x[NN:00] registers.
29	DirectMapEn: direct map enable. Read-write. 1=Enable link phy address direct map mode. This bit should only be set to access direct map link phy address registers as indicated by the register suffix ending with “_dm”.
28:16	Reserved.
15:10	UpperLinkPhyOffset: upper link phy offset address bits. Read-write.
9:0	LinkPhyOffset: link phy offset. Read-write.

D18F4x1[9C,94,8C,84] Link Phy Data Port

See [D18F4x1\[98,90,88,80\]](#) for register access information.

D18F4x1[9C,94,8C,84]_x[D0,C0] Link Phy ImpedanceTable 223: Register Mapping for [D18F4x1\[9C,94,8C,84\]_x\[D0,C0\]](#)

Register	Function
D18F4x1[9C,94,8C,84]_xC0	CAD[7:0], CTL0, CLK0
D18F4x1[9C,94,8C,84]_xD0	CAD[15:8], CTL1, CLK1

These register bits are updated as specified by [D18F0x16C\[ImmUpdate\]](#). Updates to these registers that result in a change to impedance may not take effect in the phy for up to 2 us after the update to this register completes (or until a disconnect if ImmUpdate is clear).

Bits	Description																
31:29	<p>RttCtl: receiver termination resistance (Rtt) control. Read-write. Cold reset: 0. This field specifies how the receiver termination resistance value is calculated. All values between 00h and 1Fh are valid.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Rtt is as determined by the compensation circuit, D18F4x1[9C,94,8C,84]_xE0[RttRawCal].</td> </tr> <tr> <td>001b</td> <td>Rtt is (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>101b</td> <td>If RttIndex[4]==1, Rtt is specified by the difference: RttRawCal-RttIndex[3:0]; If RttIndex[4]==0, Rtt is specified by the sum: RttRawCal + RttIndex[3:0]</td> </tr> <tr> <td>111b-110b</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. See 2.12.2 [Termination and Compensation] for more information about compensation. If RttCtl is programmed to either 011b or 100b, the value of RttRawCal + RttIndex must be less than or equal to 24.</p>	Bits	Description	000b	Rtt is as determined by the compensation circuit, D18F4x1[9C,94,8C,84]_xE0[RttRawCal] .	001b	Rtt is (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.	101b	If RttIndex[4]==1, Rtt is specified by the difference: RttRawCal-RttIndex[3:0]; If RttIndex[4]==0, Rtt is specified by the sum: RttRawCal + RttIndex[3:0]	111b-110b	Reserved
Bits	Description																
000b	Rtt is as determined by the compensation circuit, D18F4x1[9C,94,8C,84]_xE0[RttRawCal] .																
001b	Rtt is (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.																
101b	If RttIndex[4]==1, Rtt is specified by the difference: RttRawCal-RttIndex[3:0]; If RttIndex[4]==0, Rtt is specified by the sum: RttRawCal + RttIndex[3:0]																
111b-110b	Reserved																
28:21	Reserved.																
20:16	RttIndex: receiver termination resistance (Rtt) index. Read-write. Cold reset: 0. See RttCtl.																

15:13	<p>RonCtl: transmitter resistance (Ron) control. Read-write. Cold reset: 0. This field specifies how the transmitter resistance value is calculated.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Ron is as determined by the compensation circuit, D18F4x1[9C,94,8C,84]_xE0[RonRawCal].</td> </tr> <tr> <td>001b</td> <td>Ron is (RonIndex - 3).</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. See 2.12.2 [Termination and Compensation] for more information about compensation.</p> <p>If RonCtl is programmed to either 011b or 100b, the value of RonRawCal + RonIndex must be less than or equal to 23.</p>	Bits	Description	000b	Ron is as determined by the compensation circuit, D18F4x1[9C,94,8C,84]_xE0[RonRawCal] .	001b	Ron is (RonIndex - 3).	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	Description														
000b	Ron is as determined by the compensation circuit, D18F4x1[9C,94,8C,84]_xE0[RonRawCal] .														
001b	Ron is (RonIndex - 3).														
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	<p>RonIndex: transmitter resistance (Ron) index. Read-write. Cold reset: 0. See RonCtl.</p>														

D18F4x1[9C,94,8C,84]_x[D1,C1] Link Phy Receiver Loop Filter

Table 224: Register Mapping for [D18F4x1\[9C,94,8C,84\]_x\[D1,C1\]](#)

Register	Function
D18F4x1[9C,94,8C,84]_xC1	CAD[7:0], CTL0, CLK0
D18F4x1[9C,94,8C,84]_xD1	CAD[15:8], CTL1, CLK1

These register bits are updated as specified by [D18F0x16C\[ImmUpdate\]](#).




Figure 19: Link 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.

ble. The DLF counts the number of times the lane transitions before the recovered clock versus after to determine whether to adjust the recovered clock phase. 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.

Bits	Description
31:30	Reserved.
29:22	LfcMax: loop filter counter maximum value. Read-write. Cold reset: 20h.
21:14	LfcMin: loop filter counter minimum value. Read-write. Cold reset: 10h.
13:8	Reserved.
7	SamClkPiOffsetEn: sample clock phase interpolator offset enable. Read-write. Cold reset: 0. 1=Enable offset insertion around the nominal sample clock position.
6:4	SamClkPiOffset: sample clock phase interpolator offset setting. Read-write. Cold reset: X. This field specifies the magnitude of the offset of the sample clock from the nominal position. See Figure 19 . 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. Cold 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 19 .
2:0	Reserved.

D18F4x1[9C,94,8C,84]_x[D3,C3] Link Phy Timing Margin

Table 225: Register Mapping for D18F4x1[9C,94,8C,84]_x[D3,C3]

Register	Function
D18F4x1[9C,94,8C,84]_xC3	CAD[7:0], CTL0, CLK0
D18F4x1[9C,94,8C,84]_xD3	CAD[15:8], CTL1, CLK1

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.

Bits	Description
31	Reserved.
30	JitterInjEn: jitter injection enable. Read-write. Cold reset: 0. 1=Jitter injection test mode is enabled.
29	JitterInjDir: jitter injection direction. Read-write. Cold reset: 0. 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. Cold reset: 0.
22:16	Reserved.
15:10	JitterInjOffCnt:jitter injection off count. Read-write. Cold 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. Cold reset: 0. 1=Jitter injection hold is enabled.
8:0	Reserved.

D18F4x1[9C,94,8C,84]_x[D4,C4] Link Phy DFE and DFR Control

Table 226: Register Mapping for D18F4x1[9C,94,8C,84]_x[D4,C4]

Register	Function
D18F4x1[9C,94,8C,84]_xC4	CAD[7:0], CTL0, CLK0
D18F4x1[9C,94,8C,84]_xD4	CAD[15:8], CTL1, CLK1

If there is a requirement to optimize performance on a per-lane basis, see [D18F4x1\[9C,94,8C,84\]_x\[5:4\]\[9:0\]\[8,0\]6_dm\[1\]](#). These register bits are updated as specified by [D18F0x16C\[ImmUpdate\]](#).

The processor supports decision feedback restore (DFR), a function that enables on-chip AC coupling on the receiver path in Gen3 DC-coupled mode, 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. BIOS should initialize DCV as follows.

Table 227: BIOS Recommendations for D18F4x1[9C,94,8C,84]_x[D4,C4][DCV]

Far-device deemphasis setting	DCV
No deemphasis	4Dh
-2dB postcursor	3Dh
-3dB postcursor	36h
-5dB postcursor	2Bh
-6dB postcursor	27h
-7dB postcursor	22h
-8dB postcursor	1Fh
-9dB postcursor	1Bh
-11dB postcursor	16h

Decision feedback equalization (DFE) can be enabled to enhance Gen3 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.

Bits	Description
31:18	Reserved.
17:10	DCV: Vdc DAC code. Read-write. Cold reset: 0. BIOS: See Table 227 . Specifies the peak single-ended DC voltage level over the single-ended DC common mode voltage level, full swing or deemphasized, of the transmitter. DCV must be <= 87h. The expected step size is 3.125mV.
9:8	Reserved.
7	DfeEn: DFE enable. Read-write. Cold reset: 0. 1=Decision feedback equalization is enabled.

6:5	DfeVoltage: DFE offset voltage level. Read-write. Cold reset: 0. Specifies the magnitude of the DFE offset voltage. <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>25 mV</td> </tr> <tr> <td>11b-01b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	25 mV	11b-01b	Reserved
<u>Bits</u>	<u>Description</u>						
00b	25 mV						
11b-01b	Reserved						
4:0	Reserved.						

D18F4x1[9C,94,8C,84]_x[D5,C5] Link Phy Transmit Control

Table 228: **Register Mapping** for D18F4x1[9C,94,8C,84]_x[D5,C5]

Register	Function
D18F4x1[9C,94,8C,84]_xC5	CAD[7:0], CTL0, CLK0
D18F4x1[9C,94,8C,84]_xD5	CAD[15:8], CTL1, CLK1

Bits	Description
31:8	Reserved.
7	TxLs23ClkGateEn: LS2/LS3 clock gating enable. Read-write. Cold reset: 1. 1=Internal phy clock grids are gated during LS2 or PHY OFF states to save power. Recommended to set TxLs23ClkGateEn for systems that support LDTSTOP and LS2.
6:0	Reserved.

D18F4x1[9C,94,8C,84]_x[DF,CF] Link FIFO Read Pointer Optimization

Table 229: Register Mapping for D18F4x1[9C,94,8C,84]_x[DF,CF]

Register	Function
D18F4x1[9C,94,8C,84]_xCF	CAD[7:0], CTL0, CLK0
D18F4x1[9C,94,8C,84]_xDF	CAD[15:8], CTL1, CLK1

There is a synchronization FIFO between the NB clock domain and each of the link clock domains. At cold reset, the read pointer and write pointer for each of these FIFOs is positioned conservatively (30 bit-times apart), such that FIFO latency may be greater than is necessary. This register may be used to position the read pointer and write pointer of each FIFO closer to each other such that latency is reduced. Each of the fields of this register specify the number of positions to move read pointer closer to the write pointer.

If the pointers were at minimal separation during normal operation, a reduction in VDDNB while the transmitter is running would result in a FIFO underrun as the write pointer slows down. Hardware can automatically offset the pointers according to the values specified in this register to compensate for the VDDNB changes to prevent the FIFO underrun issue. This register provides the capability to uniquely program the pointer separation for each of the 4 PStates; P0XmtRdPtr and P0RcvRdPtr for P0 P-state which corresponds to the highest VDDNB power state for maximum performance, P3XmtRdPtr and P3RcvRdPtr for P3 P-state which corresponds to the lowest VDDNB power state, and P1 and P2 for the 2 intermediate PStates.

After writing to this register, the new values are applied to the FIFOs each time the link disconnects and reconnects, including warm resets and LDTSTOP_L assertions. If there is a transition to a higher P-state, the Rcv and Xmt pointer adjust register settings for the higher P-state to be transitioned into is applied. If there is a transition to a lower P-state or no transition, the Rcv and Xmt pointer adjust register settings for the P0 P-state is applied. Typically, the RCV FIFO Async tracking logic can compensate for any decrease in VDDNB during normal operation; hence, all 4 of the RcvRdPtr settings are expected to be the same. Async clocking mode does not move the pointers closer than programmed, it only allows them to keep the programmed separation when the received clock is faster or slower than the transmit clock.

Reads from P0XmtRdPtr and P0RcvRdPtr after a write but before the link disconnects and reconnects, return the current value, not the pending value from the last write. Reads from P1XmtRdPtr, P1RcvRdPtr, P2XmtRdPtr, P2RcvRdPtr, P3XmtRdPtr and P3RcvRdPtr always return 0h.

Table 230: BIOS Recommendations for P0RcvRdPtr, P1RcvRdPtr.

Condition	D18F4x1[9C,94,8C,84]_x[DF,CF]	
HTCLK (GHz)	P0RcvRdPtr	P1RcvRdPtr
~Gen3	Dh	Dh
Gen3	Ah	Ah

Table 231: BIOS Recommendations for P0XmtRdPtr, P1XmtRdPtr

Condition	D18F4x1[9C,94,8C,84]_x[DF,CF]	
HTCLK (GHz)	P0XmtRdPtr	P1XmtRdPtr
~Gen3	4	0
Gen3	2	0

Bits	Description										
31:28	P3XmtRdPtr: transmit FIFO read pointer for NB P3. RAZ; write; reset-applied. Cold reset: 0. BIOS: 0. See P0XmtRdPtr.										
27:24	P3RcvRdPtr: receive FIFO read pointer for NB P3. RAZ; write; reset-applied. Cold reset: 0. BIOS: 0. See P0RcvRdPtr.										
23:20	P2XmtRdPtr: transmit FIFO read pointer for NB P2. RAZ; write; reset-applied. Cold reset: 0. BIOS: 0. See P0XmtRdPtr.										
19:16	P2RcvRdPtr: receive FIFO read pointer for NB P2. RAZ; write; reset-applied. Cold reset: 0. BIOS: 0. See P0RcvRdPtr.										
15:12	P1XmtRdPtr: transmit FIFO read pointer for NB P1. RAZ; write; reset-applied. Cold reset: 0. BIOS: Table 231 . See P0XmtRdPtr.										
11:8	P1RcvRdPtr: receive FIFO read pointer for NB P1. RAZ; write; reset-applied. Cold reset: 0. BIOS: Table 230 . See P0RcvRdPtr.										
7:4	<p>P0XmtRdPtr: transmit FIFO read pointer for NB P0. Read-write; reset-applied. Cold reset: 0. BIOS: Table 231. Specified in double-bit time increments. Reads returns the current value, not the pending value from the last write.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Position the read pointer 0 bit times closer to the write pointer.</td> </tr> <tr> <td>1h</td> <td>Position the read pointer 2 bit times closer to the write pointer.</td> </tr> <tr> <td>Eh-2h</td> <td>Position the read pointer <P0XmtRdPtr*2> bit times closer to the write pointer.</td> </tr> <tr> <td>Fh</td> <td>Position the read pointer 30 bit times closer to the write pointer.</td> </tr> </tbody> </table>	Bits	Description	0h	Position the read pointer 0 bit times closer to the write pointer.	1h	Position the read pointer 2 bit times closer to the write pointer.	Eh-2h	Position the read pointer <P0XmtRdPtr*2> bit times closer to the write pointer.	Fh	Position the read pointer 30 bit times closer to the write pointer.
Bits	Description										
0h	Position the read pointer 0 bit times closer to the write pointer.										
1h	Position the read pointer 2 bit times closer to the write pointer.										
Eh-2h	Position the read pointer <P0XmtRdPtr*2> bit times closer to the write pointer.										
Fh	Position the read pointer 30 bit times closer to the write pointer.										
3:0	<p>P0RcvRdPtr: receive FIFO read pointer for NB P0. Read-write; reset-applied. Cold reset: 0. BIOS: See Table 230. Specified in double-bit time increments. Reads returns the current value, not the pending value from the last write.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Position the read pointer 0 bit times closer to the write pointer.</td> </tr> <tr> <td>1h</td> <td>Position the read pointer 2 bit times closer to the write pointer.</td> </tr> <tr> <td>Eh-2h</td> <td>Position the read pointer <P0RcvRdPtr*2> bit times closer to the write pointer.</td> </tr> <tr> <td>Fh</td> <td>Position the read pointer 30 bit times closer to the write pointer.</td> </tr> </tbody> </table>	Bits	Description	0h	Position the read pointer 0 bit times closer to the write pointer.	1h	Position the read pointer 2 bit times closer to the write pointer.	Eh-2h	Position the read pointer <P0RcvRdPtr*2> bit times closer to the write pointer.	Fh	Position the read pointer 30 bit times closer to the write pointer.
Bits	Description										
0h	Position the read pointer 0 bit times closer to the write pointer.										
1h	Position the read pointer 2 bit times closer to the write pointer.										
Eh-2h	Position the read pointer <P0RcvRdPtr*2> bit times closer to the write pointer.										
Fh	Position the read pointer 30 bit times closer to the write pointer.										

D18F4x1[9C,94,8C,84]_xE0 Link Phy Compensation and Calibration Control 1

These register bits are updated as specified by [D18F0x16C](#)[ImmUpdate].

Bits	Description
31:28	Reserved.
27:23	RttRawCal: receiver termination resistance (Rtt) raw calibration value. Read-only. Cold 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. Cold reset: 0. This field provides the raw Ron calibration value as determined by the compensation circuit.
17:0	Reserved.

D18F4x1[9C,94,8C,84]_xE3 Link Phy PLL Control

These register bits are updated as specified by [D18F0x16C](#)[ImmUpdate].

Bits	Description
31:8	Reserved.
7	RoCalEn: RO calibration enable. Read-write. Cold reset: 0. BIOS: 1. 1=Enables RO calibration.
6:0	Reserved.

D18F4x1[9C,94,8C,84]_x100 Link BIST Control

Bits	Description												
31	Width. Read-only. Indicates the implemented width of the BIST engine. 0=8 bits. In 16-bit links, the same patterns are transmitted on the upper and lower sublinks. The upper bit of D18F0x[18C:170][LaneSel] selects which half of the link is checked in the receiver.												
30:27	Reserved.												
26:16	ErrCnt: error count. Read; write-1-to-clear; updated-by-hardware. Cold reset: 0. This field is incremented by hardware upon detection of each error on any lane. This count is the sum of error counts from each lane, each of which saturates at 63. Writes other than all-zeroes or all-ones result in undefined behavior.												
15:14	Reserved.												
13	LinkInitFailure: link initialization failure. Read-write; set-by-hardware; write-1-to-clear. Cold-reset: 0. 1=Timeout in TR1/TR2 or exit through TM4.												
12:8	ErrLnNum: error lane number. Read; write-1-to-clear; updated-by-hardware. Cold reset: 0. This value is set by hardware to the lane of the sublink that failed upon detection of the first error by the BIST receiver. If multiple bits fail at the same time, the highest-numbered bit is recorded. Writes other than all-zeroes or all-ones result in undefined behavior. <table border="1" data-bbox="267 1134 747 1344"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>CAD0</td> </tr> <tr> <td>6h-1h</td> <td>CAD<ErrLnNum></td> </tr> <tr> <td>7h</td> <td>CAD7</td> </tr> <tr> <td>8h</td> <td>CTL</td> </tr> <tr> <td>Fh-9h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	CAD0	6h-1h	CAD<ErrLnNum>	7h	CAD7	8h	CTL	Fh-9h	Reserved
Bits	Description												
0h	CAD0												
6h-1h	CAD<ErrLnNum>												
7h	CAD7												
8h	CTL												
Fh-9h	Reserved												
7:6	ErrStat: error status. Read; write-1-to-clear; updated-by-hardware. Cold reset: 00b. This value is set by hardware to the error type upon detection of the first error by the BIST receiver. Writes other than all-zeroes or all-ones result in undefined behavior. <table border="1" data-bbox="267 1449 747 1627"> <thead> <tr> <th>Bits</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>No error</td> </tr> <tr> <td>01b</td> <td>In training</td> </tr> <tr> <td>10b</td> <td>Pattern miscompare</td> </tr> <tr> <td>11b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Status	00b	No error	01b	In training	10b	Pattern miscompare	11b	Reserved		
Bits	Status												
00b	No error												
01b	In training												
10b	Pattern miscompare												
11b	Reserved												
5	InvRotEn: inversion rotate enable. Read-write. Cold reset: 0. This bit enables rotation of D18F4x1[9C,94,8C,84]_x110 [Link BIST Southbound TX Inversion] and D18F4x1[9C,94,8C,84]_x130 [Link BIST Northbound RX Inversion] at the completion of each BIST loop.												
4:2	Reserved.												

1	RxDis: receiver disable. Read-write. Cold reset: 0. 1=Disables checking of BIST patterns in the receiver if BIST is already active. An LDTSTOP_L or RESET_L assertion is still required to exit BIST. If BIST has not started yet, setting this bit additionally removes any dependency on receiver link training, such that the transmitter sequences through the minimum training sets and begin sending BIST patterns at the completion of these training sets.
0	Reserved.

D18F4x1[9C,94,8C,84]_x104 Link BIST Southbound TX Pattern Control

Bits	Description																		
31:26	Reserved.																		
25:21	ConstCnt: constant generator count. Read-write. Cold reset: 0. Selects the number of times to repeat the constant selected by ConstSel, in multiples of 24 bits. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>0 (disabled)</td> </tr> <tr> <td>01h</td> <td>24 bits</td> </tr> <tr> <td>1Eh-02h</td> <td><ConstCnt*24> bits</td> </tr> <tr> <td>1Fh</td> <td>744 bits</td> </tr> </tbody> </table>	Bits	Description	00h	0 (disabled)	01h	24 bits	1Eh-02h	<ConstCnt*24> bits	1Fh	744 bits								
Bits	Description																		
00h	0 (disabled)																		
01h	24 bits																		
1Eh-02h	<ConstCnt*24> bits																		
1Fh	744 bits																		
20	ConstSel: constant generator select. Read-write. Cold reset: 0. Selects 0 or 1 to send for the time the constant generator is active.																		
19:13	ModCnt: modulo-N count. Read-write. Cold reset: 0. Selects the number of times to repeat the Modulo-N counter (a counter with a period of N bits) pattern, 0 to 127.																		
12:10	ModSel: modulo-N select. Read-write. Cold reset: 000b. Selects the pattern sent by the Modulo-N counter. <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>L/2 – 0101_0101_0101_0101_0101_0101b</td> </tr> <tr> <td>010b</td> <td>L/4 – 0011_0011_0011_0011_0011_0011b</td> </tr> <tr> <td>011b</td> <td>L/6 – 0001_1100_0111_0001_1100_0111b</td> </tr> <tr> <td>100b</td> <td>L/8 – 0000_1111_0000_1111_0000_1111b</td> </tr> <tr> <td>101b</td> <td>Reserved</td> </tr> <tr> <td>110b</td> <td>L/24 – 0000_0000_0000_1111_1111_1111b</td> </tr> <tr> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	Reserved	001b	L/2 – 0101_0101_0101_0101_0101_0101b	010b	L/4 – 0011_0011_0011_0011_0011_0011b	011b	L/6 – 0001_1100_0111_0001_1100_0111b	100b	L/8 – 0000_1111_0000_1111_0000_1111b	101b	Reserved	110b	L/24 – 0000_0000_0000_1111_1111_1111b	111b	Reserved
Bits	Description																		
000b	Reserved																		
001b	L/2 – 0101_0101_0101_0101_0101_0101b																		
010b	L/4 – 0011_0011_0011_0011_0011_0011b																		
011b	L/6 – 0001_1100_0111_0001_1100_0111b																		
100b	L/8 – 0000_1111_0000_1111_0000_1111b																		
101b	Reserved																		
110b	L/24 – 0000_0000_0000_1111_1111_1111b																		
111b	Reserved																		
9:3	PatCnt: pattern buffer count. Read-write. Cold reset: 0. Selects the number of times to repeat the pattern selected by D18F4x1[9C,94,8C,84]_x118, 0 to 127.																		
2:0	Order. Read-write. Cold reset: 0. Selects the order in which each pattern is sent. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Pattern Buffer, Modulo-N Counter, Constant Generator</td> </tr> <tr> <td>001b</td> <td>Pattern Buffer, Constant Generator, Modulo-N Counter</td> </tr> <tr> <td>010b</td> <td>Modulo-N Counter, Pattern Buffer, Constant Generator</td> </tr> <tr> <td>011b</td> <td>Modulo-N Counter, Constant Generator, Pattern Buffer</td> </tr> <tr> <td>100b</td> <td>Constant Generator, Pattern Buffer, Modulo-N Counter</td> </tr> <tr> <td>101b</td> <td>Constant Generator, Modulo-N Counter, Pattern Buffer</td> </tr> <tr> <td>11xb</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	000b	Pattern Buffer, Modulo-N Counter, Constant Generator	001b	Pattern Buffer, Constant Generator, Modulo-N Counter	010b	Modulo-N Counter, Pattern Buffer, Constant Generator	011b	Modulo-N Counter, Constant Generator, Pattern Buffer	100b	Constant Generator, Pattern Buffer, Modulo-N Counter	101b	Constant Generator, Modulo-N Counter, Pattern Buffer	11xb	Reserved		
Bits	Description																		
000b	Pattern Buffer, Modulo-N Counter, Constant Generator																		
001b	Pattern Buffer, Constant Generator, Modulo-N Counter																		
010b	Modulo-N Counter, Pattern Buffer, Constant Generator																		
011b	Modulo-N Counter, Constant Generator, Pattern Buffer																		
100b	Constant Generator, Pattern Buffer, Modulo-N Counter																		
101b	Constant Generator, Modulo-N Counter, Pattern Buffer																		
11xb	Reserved																		

D18F4x1[9C,94,8C,84]_x108 Link BIST Southbound TX Pattern Buffer 1

Bits	Description
31:24	Reserved.
23:0	Pattern1[23:0] . Read-write. Cold reset: 0. Holds the first 24 bits of Pattern Buffer 1.

D18F4x1[9C,94,8C,84]_x10C Link BIST Southbound TX Mask

Bits	Description										
31:9	Reserved.										
8:0	<p>TxMask. Read-write. Cold reset: 1FFh. Selects lanes of the sublinks to transmit a logical 0. 1=Lane active. 0=Lane masked.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<TxMask></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table>	Bit	Description	[0]	CAD0	[6:1]	CAD<TxMask>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<TxMask>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x110 Link BIST Southbound TX Inversion

Bits	Description										
31:9	Reserved.										
8:0	<p>TxInv. Read-write. Cold reset: 0. Selects lanes of the sublinks to invert. 1=Lane inverted. 0=Lane unmodified.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<TxInv></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table> <p>When D18F4x1[9C,94,8C,84]_x100[InvRotEn] is set, the bits corresponding to active lanes rotate to the left at the completion of each BIST loop: {NxtTxInv[8:0]}={TxInv[7:0],TxInv[8]}. If the transmitter and receiver are different widths, inversion rotation can only be used for 16/8-bit links and the initial pattern in the inversion register must repeat on 9-bit boundaries.</p>	Bit	Description	[0]	CAD0	[6:1]	CAD<TxInv>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<TxInv>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x114 Link BIST Southbound TX Pattern Buffer 2

Bits	Description
31:24	Reserved.
23:0	Pattern2[23:0] . Read-write. Cold reset: 0. Holds the first 24 bits of Pattern Buffer 2.

D18F4x1[9C,94,8C,84]_x118 Link BIST Southbound TX Pattern Buffer 2 Enable

Bits	Description										
31:9	Reserved.										
8:0	Pat2En. Read-write. Cold reset: 0. Selects lanes of the sublinks that use Pattern Buffer 2 instead of Pattern Buffer 1. 1=Buffer 2 selected. 0=Buffer 1 selected. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<Pat2En></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table>	Bit	Description	[0]	CAD0	[6:1]	CAD<Pat2En>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<Pat2En>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x11C Link BIST Southbound TX Pattern Buffer Extension

Bits	Description
31:16	Pattern2[39:24]. Read-write. Cold reset: 0. Holds the upper 16 bits of Pattern Buffer 2.
15:0	Pattern1[39:24]. Read-write. Cold reset: 0. Holds the upper 16 bits of Pattern Buffer 1.

D18F4x1[9C,94,8C,84]_x120 Link BIST Southbound TX Scramble

Bits	Description										
31:9	Reserved.										
8:0	TxScramble. Read-write. Cold reset: 0. Selects lanes of the sublinks to scramble. 1=Scrambling enabled. 0=Scrambling disabled. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<TxScramble></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table>	Bit	Description	[0]	CAD0	[6:1]	CAD<TxScramble>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<TxScramble>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x124 Link BIST Northbound RX Pattern Control

Bits	Description						
31:26	Reserved.						
25:21	ConstCnt: constant generator count. Read-write. Cold reset: 0. Selects the number of times to repeat the constant selected by ConstSel, in multiples of 24 bits. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00000b</td> <td>0 (disabled)</td> </tr> <tr> <td>11111b-00001b</td> <td>24*ConstCnt bits (24*1=24 to 24*31=744 bits)</td> </tr> </tbody> </table>	Bits	Description	00000b	0 (disabled)	11111b-00001b	24*ConstCnt bits (24*1=24 to 24*31=744 bits)
Bits	Description						
00000b	0 (disabled)						
11111b-00001b	24*ConstCnt bits (24*1=24 to 24*31=744 bits)						
20	ConstSel: constant generator select. Read-write. Cold reset: 0. Selects 0 or 1 to send for the time the constant generator is active.						
19:13	ModCnt: modulo-N count. Read-write. Cold reset: 0. Selects the number of times to repeat the Modulo-N counter (a counter with a period of N bits) pattern, 0 to 127.						

12:10	ModSel: modulo-N select. Read-write. Cold reset: 000b. Selects the pattern sent by the Modulo-N counter. <u>Bits</u> 000b 001b 010b 011b 100b 101b 110b 111b	<u>Description</u> Reserved L/2 – 0101_0101_0101_0101_0101b L/4 – 0011_0011_0011_0011_0011b L/6 – 0001_1100_0111_0001_1100_0111b L/8 – 0000_1111_0000_1111_0000_1111b Reserved L/24 – 0000_0000_0000_1111_1111_1111b Reserved
9:3	PatCnt: pattern buffer count. Read-write. Cold reset: 0. Selects the number of times to repeat the pattern selected by D18F4x1[9C,94,8C,84]_x118 , 0 to 127.	
2:0	Order. Read-write. Cold reset: 0. Selects the order in which each pattern is sent. <u>Bits</u> 000b 001b 010b 011b 100b 101b 111b-110b	<u>Description</u> Pattern Buffer, Modulo-N Counter, Constant Generator Pattern Buffer, Constant Generator, Modulo-N Counter Modulo-N Counter, Pattern Buffer, Constant Generator Modulo-N Counter, Constant Generator, Pattern Buffer Constant Generator, Pattern Buffer, Modulo-N Counter Constant Generator, Modulo-N Counter, Pattern Buffer Reserved

D18F4x1[9C,94,8C,84]_x128 Link BIST Northbound RX Pattern Buffer 1

Bits	Description
31:24	Reserved.
23:0	Pattern1[23:0]. Read-write. Cold reset: 0. Holds the first 24 bits of Pattern Buffer 1.

D18F4x1[9C,94,8C,84]_x12C Link BIST Northbound RX Mask

Bits	Description
31:9	Reserved.
8:0	RxMask. Read-write. Cold reset: 1FFh. Selects lanes of the selected sublink that are checked by the receiver. 1=Lane active. 0=Lane masked. Software is responsible for clearing bits 7:4 for 4-bit links and bits 7:2 for 2-bit links. <u>Bit</u> [0] [6:1] [7] [8]
	<u>Description</u> CAD0 CAD<RxMask> CAD7 CTL

D18F4x1[9C,94,8C,84]_x130 Link BIST Northbound RX Inversion

Bits	Description										
31:9	Reserved.										
8:0	<p>RxInv. Read-write. Cold reset: 0. Selects lanes of the sublink that are inverted. 1=Lane inverted. 0=Lane unmodified.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<RxInv></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table> <p>When D18F4x1[9C,94,8C,84]_x100[InvRotEn] is set, the bits corresponding to active lanes rotate to the left at the completion of each BIST loop: {NxtTxInv[8:0]}={TxInv[7:0],TxInv[8]}. If the transmitter and receiver are different widths, inversion rotation can only be used for 16/8-bit links and the initial pattern in the inversion register must repeat on 9-bit boundaries.</p>	Bit	Description	[0]	CAD0	[6:1]	CAD<RxInv>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<RxInv>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x134 Link BIST Northbound RX Pattern Buffer 2

Bits	Description
31:24	Reserved.
23:0	Pattern2[23:0] . Read-write. Cold reset: 0. Holds the first 24 bits of Pattern Buffer 2.

D18F4x1[9C,94,8C,84]_x138 Link BIST Northbound RX Pattern Buffer 2 Enable

Bits	Description										
31:9	Reserved.										
8:0	<p>Pat2En. Read-write. Cold reset: 0. Selects lanes of the sublink that use Pattern Buffer 2 instead of Pattern Buffer 1. 1=Buffer 2 selected. 0=Buffer 1 selected.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<Pat2En></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table>	Bit	Description	[0]	CAD0	[6:1]	CAD<Pat2En>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<Pat2En>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x13C Link BIST Northbound RX Pattern Buffer Extension

Bits	Description
31:16	Pattern2[39:24] . Read-write. Cold reset: 0. Holds the upper 16 bits of Pattern Buffer 2.
15:0	Pattern1[39:24] . Read-write. Cold reset: 0. Holds the upper 16 bits of Pattern Buffer 1.

D18F4x1[9C,94,8C,84]_x140 Link BIST Northbound RX Scramble

Bits	Description										
31:9	Reserved.										
8:0	RxScramble. Read-write. Cold reset: 0. Selects lanes of the sublink to scramble. 1=Scrambling enabled. 0=Scrambling disabled. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<RxScramble></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table>	Bit	Description	[0]	CAD0	[6:1]	CAD<RxScramble>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<RxScramble>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x144 Link BIST Northbound RX Error Status

Bits	Description										
31	Reserved.										
30:24	ErrCntCtl: CTL lane Error Count. Read-only. ErrCntCtl[5:0] indicates the number of BIST errors detected on this lane since the last time D18F4x1[9C,94,8C,84]_x100[ErrCnt] was cleared. ErrCntCtl[6] indicates overflow when set.										
23:9	Reserved.										
8:0	RxErrStat. Read; write 000h to clear (if non-0 then the write is ignored); set-by-hardware. Cold reset: 0. Indicates lanes of the selected sublink that had errors. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>CAD0</td> </tr> <tr> <td>[6:1]</td> <td>CAD<RxErrStat></td> </tr> <tr> <td>[7]</td> <td>CAD7</td> </tr> <tr> <td>[8]</td> <td>CTL</td> </tr> </tbody> </table>	Bit	Description	[0]	CAD0	[6:1]	CAD<RxErrStat>	[7]	CAD7	[8]	CTL
Bit	Description										
[0]	CAD0										
[6:1]	CAD<RxErrStat>										
[7]	CAD7										
[8]	CTL										

D18F4x1[9C,94,8C,84]_x148 Link BIST Northbound RX Per-Lane Error Count 1

Bits	Description
31	Reserved.
30:24	ErrCntCad3: CAD3 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .
23	Reserved.
22:16	ErrCntCad2: CAD2 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .
15	Reserved.
14:8	ErrCntCad1: CAD1 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .
7	Reserved.
6:0	ErrCntCad0: CAD0 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .

D18F4x1[9C,94,8C,84]_x14C Link BIST Northbound RX Per-Lane Error Count 2

Bits	Description
31	Reserved.
30:24	ErrCntCad7: CAD7 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .
23	Reserved.
22:16	ErrCntCad6: CAD6 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .
15	Reserved.
14:8	ErrCntCad5: CAD5 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .
7	Reserved.
6:0	ErrCntCad4: CAD4 lane Error Count. See: D18F4x1[9C,94,8C,84]_x144[ErrCntCtl] .

D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]6_dm[1] Link Phy DFE and DFR Control

This is a direct map register set; see [D18F4x1\[98,90,88,80\]\[DirectMapEn\]](#). The definition of this register is similar to that of [D18F4x1\[9C,94,8C,84\]_x\[D4,C4\]](#). This register provides per-lane programmability whereas [D18F4x1\[9C,94,8C,84\]_x\[D4,C4\]](#) only allows per-sublink wide programmability.

Table 232: Register Mapping for [D18F4x1\[9C,94,8C,84\]_x\[5:4\]\[9:0\]\[8,0\]6_dm\[1\]](#)

Register	Function
D18F4x1[9C,94,8C,84]_x4006_dm1	CADIN[0]
D18F4x1[9C,94,8C,84]_x4086_dm1	CADIN[1]
D18F4x1[9C,94,8C,84]_x4106_dm1	CADIN[2]
D18F4x1[9C,94,8C,84]_x4186_dm1	CADIN[3]
D18F4x1[9C,94,8C,84]_x4206_dm1	CADIN[4]
D18F4x1[9C,94,8C,84]_x4286_dm1	CADIN[5]
D18F4x1[9C,94,8C,84]_x4306_dm1	CADIN[6]
D18F4x1[9C,94,8C,84]_x4386_dm1	CADIN[7]
D18F4x1[9C,94,8C,84]_x4406_dm1	CADIN[8]
D18F4x1[9C,94,8C,84]_x4486_dm1	CADIN[9]
D18F4x1[9C,94,8C,84]_x4506_dm1	CADIN[10]
D18F4x1[9C,94,8C,84]_x4586_dm1	CADIN[11]
D18F4x1[9C,94,8C,84]_x4606_dm1	CADIN[12]
D18F4x1[9C,94,8C,84]_x4686_dm1	CADIN[13]
D18F4x1[9C,94,8C,84]_x4706_dm1	CADIN[14]
D18F4x1[9C,94,8C,84]_x4786_dm1	CADIN[15]
D18F4x1[9C,94,8C,84]_x4806_dm1	CTLIN[0]
D18F4x1[9C,94,8C,84]_x4886_dm1	Reserved
D18F4x1[9C,94,8C,84]_x4906_dm1	CTLIN[1]
D18F4x1[9C,94,8C,84]_x4986_dm1	Reserved

Table 233: Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]6_dm[1]

Link width	D18F4x1[98,90,88,80][15:8]	D18F4x1[98,90,88,80][7:0]
		06h
16-bit	50h	D18F4x1[8C,94,8C,84]_x[4[7:0][8,0]6,4[9:8]06]
8-bit	53h	D18F4x1[8C,94,8C,84]_x[4[7:4][8,0]6,4906]
8-bit	52h	D18F4x1[8C,94,8C,84]_x[4[3:0][8,0]6,4806]
4-bit	57h	D18F4x1[8C,94,8C,84]_x[4[5:4][8,0]6,4906]
4-bit	56h	D18F4x1[8C,94,8C,84]_x[4[1:0][8,0]6,4806]
2-bit	55h	D18F4x1[8C,94,8C,84]_x[44[8,0]6,4906]
2-bit	54h	D18F4x1[8C,94,8C,84]_x[40[8,0]6,4806]

Accessing any register number that is not listed above might result in undefined behavior of the phy. These register bits are updated as specified by D18F0x16C[ImmUpdate].

Bits	Description						
31:18	Reserved.						
17:10	DCV: Vdc DAC code. Read-write. Cold reset: 0. BIOS: See Table 227. Specifies the peak single-ended DC voltage level over the single-ended DC common mode voltage level, full swing or deemphasized, of the transmitter. DCV must be <= 87h. The expected step size is 3.125mV.						
9:8	Reserved.						
7	DfeEn: DFE enable. Read-write. Cold reset: 0. 1=Decision feedback equalization is enabled.						
6:5	DfeVoltage: DFE offset voltage level. Read-write. Cold reset: 0. Specifies the magnitude of the DFE offset voltage. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>25 mV</td> </tr> <tr> <td>11b-01b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	00b	25 mV	11b-01b	Reserved
Bits	Description						
00b	25 mV						
11b-01b	Reserved						
4:0	Reserved.						

D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1] Link Phy DLL Control

This is a direct map register set; see D18F4x1[98,90,88,80][DirectMapEn].

Table 234: Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1]

Register	Function
D18F4x1[9C,94,8C,84]_x400A_dm1	CADIN[0]
D18F4x1[9C,94,8C,84]_x408A_dm1	CADIN[1]
D18F4x1[9C,94,8C,84]_x410A_dm1	CADIN[2]
D18F4x1[9C,94,8C,84]_x418A_dm1	CADIN[3]
D18F4x1[9C,94,8C,84]_x420A_dm1	CADIN[4]
D18F4x1[9C,94,8C,84]_x428A_dm1	CADIN[5]
D18F4x1[9C,94,8C,84]_x430A_dm1	CADIN[6]
D18F4x1[9C,94,8C,84]_x438A_dm1	CADIN[7]
D18F4x1[9C,94,8C,84]_x440A_dm1	CADIN[8]

Table 234: Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1]

D18F4x1[9C,94,8C,84]_x448A_dm1	CADIN[9]
D18F4x1[9C,94,8C,84]_x450A_dm1	CADIN[10]
D18F4x1[9C,94,8C,84]_x458A_dm1	CADIN[11]
D18F4x1[9C,94,8C,84]_x460A_dm1	CADIN[12]
D18F4x1[9C,94,8C,84]_x468A_dm1	CADIN[13]
D18F4x1[9C,94,8C,84]_x470A_dm1	CADIN[14]
D18F4x1[9C,94,8C,84]_x478A_dm1	CADIN[15]
D18F4x1[9C,94,8C,84]_x480A_dm1	CTLIN[0]
D18F4x1[9C,94,8C,84]_x488A_dm1	Reserved
D18F4x1[9C,94,8C,84]_x490A_dm1	CTLIN[1]
D18F4x1[9C,94,8C,84]_x498A_dm1	Reserved

Table 235: Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1]

Link width	D18F4x1[98,90,88,80][15:8]	D18F4x1[98,90,88,80][7:0]
		0Ah
16-bit	50h	D18F4x1[8C,94,8C,84]_x[4[7:0][8,0]A,4[9:8]0A]
8-bit	53h	D18F4x1[8C,94,8C,84]_x[4[7:4][8,0]A,490A]
8-bit	52h	D18F4x1[8C,94,8C,84]_x[4[3:0][8,0]A,480A]
4-bit	57h	D18F4x1[8C,94,8C,84]_x[4[5:4][8,0]A,490A]
4-bit	56h	D18F4x1[8C,94,8C,84]_x[4[1:0][8,0]A,480A]
2-bit	55h	D18F4x1[8C,94,8C,84]_x[44[8,0]A,490A]
2-bit	54h	D18F4x1[8C,94,8C,84]_x[40[8,0]A,480A]

Bits	Description																
31:29	<p>Ls2ExitTime: LS2 exit time. Read-write. Cold reset: 0. BIOS: 0. 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>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>10 us</td> </tr> <tr> <td>001b</td> <td>5 us</td> </tr> <tr> <td>010b</td> <td>2.5 us</td> </tr> <tr> <td>011b</td> <td>1.25 us</td> </tr> <tr> <td>100b</td> <td>625 ns</td> </tr> <tr> <td>101b</td> <td>0 s</td> </tr> <tr> <td>111b-110b</td> <td>Reserved</td> </tr> </tbody> </table> <p>The value specified by Ls2ExitTime must be less than the value specified by D18F0x16C[T0Time], or it can cause undefined behavior.</p>	Bits	Description	000b	10 us	001b	5 us	010b	2.5 us	011b	1.25 us	100b	625 ns	101b	0 s	111b-110b	Reserved
Bits	Description																
000b	10 us																
001b	5 us																
010b	2.5 us																
011b	1.25 us																
100b	625 ns																
101b	0 s																
111b-110b	Reserved																
28:15	Reserved.																

14:13	AnalogWaitTime: analog wait time to turn on DLL. Read-write. Cold reset: 0. BIOS: 10b. Specifies the delay from cold reset to turning on the DLL circuit. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>1.25 us</td> </tr> <tr> <td>01b</td> <td>0.625 us</td> </tr> <tr> <td>10b</td> <td>2.5 us</td> </tr> <tr> <td>11b</td> <td>0.3125 us</td> </tr> </tbody> </table>	Bits	Description	00b	1.25 us	01b	0.625 us	10b	2.5 us	11b	0.3125 us
Bits	Description										
00b	1.25 us										
01b	0.625 us										
10b	2.5 us										
11b	0.3125 us										
12:8	Reserved.										
7	BiasDisInLs2: bias disable in LS2 power state. Read-write. Cold reset: 0. BIOS: 0. 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: CDR lock detect on LS2 exit. Read-write. Cold reset: 0. BIOS: 0. Selects the LS2 to L0 power state transition speed. 1=Fast transition mode selected. 0=Slow transition mode selected.										
3:1	Reserved.										
0	EnCoreLoopFirst: enable DLL core loop first on LS2 exit. Read-write. Cold reset: 0. BIOS: 0. This field selects LS2 to L0 power state transition speed. 1=Fast transition mode selected. 0=Slow transition mode selected.										

D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]F_dm[1] Link Phy Receiver DLL Control and Test 5

This is a direct map register set; see [D18F4x1\[98,90,88,80\]\[DirectMapEn\]](#). These registers provide test and debug features control associated with the link receive lanes.

Table 236: [Register Mapping](#) for [D18F4x1\[9C,94,8C,84\]_x\[5:4\]\[9:0\]\[8,0\]F_dm\[1\]](#)

Register	Function
D18F4x1[9C,94,8C,84]_x400F_dm1	CADIN[0]
D18F4x1[9C,94,8C,84]_x408F_dm1	CADIN[1]
D18F4x1[9C,94,8C,84]_x410F_dm1	CADIN[2]
D18F4x1[9C,94,8C,84]_x418F_dm1	CADIN[3]
D18F4x1[9C,94,8C,84]_x420F_dm1	CADIN[4]
D18F4x1[9C,94,8C,84]_x428F_dm1	CADIN[5]
D18F4x1[9C,94,8C,84]_x430F_dm1	CADIN[6]
D18F4x1[9C,94,8C,84]_x438F_dm1	CADIN[7]
D18F4x1[9C,94,8C,84]_x440F_dm1	CADIN[8]
D18F4x1[9C,94,8C,84]_x448F_dm1	CADIN[9]
D18F4x1[9C,94,8C,84]_x450F_dm1	CADIN[10]
D18F4x1[9C,94,8C,84]_x458F_dm1	CADIN[11]
D18F4x1[9C,94,8C,84]_x460F_dm1	CADIN[12]
D18F4x1[9C,94,8C,84]_x468F_dm1	CADIN[13]
D18F4x1[9C,94,8C,84]_x470F_dm1	CADIN[14]

Table 236: Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]F_dm[1]

D18F4x1[9C,94,8C,84]_x478F_dm1	CADIN[15]
D18F4x1[9C,94,8C,84]_x480F_dm1	CTLIN[0]
D18F4x1[9C,94,8C,84]_x490F_dm1	CTLIN[1]

Table 237: Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]F_dm[1]

Link width	D18F4x1[98,90,88,80][15:8]	D18F4x1[98,90,88,80][7:0]
		0Fh
16-bit	50h	D18F4x1[8C,94,8C,84]_x[4[7:0][8,0]F,4[9:8]0F]
8-bit	53h	D18F4x1[8C,94,8C,84]_x[4[7:4][8,0]F,490F]
8-bit	52h	D18F4x1[8C,94,8C,84]_x[4[3:0][8,0]F,480F]
4-bit	57h	D18F4x1[8C,94,8C,84]_x[4[5:4][8,0]F,490F]
4-bit	56h	D18F4x1[8C,94,8C,84]_x[4[1:0][8,0]F,480F]
2-bit	55h	D18F4x1[8C,94,8C,84]_x[44[8,0]F,490F]
2-bit	54h	D18F4x1[8C,94,8C,84]_x[40[8,0]F,480F]

Accessing any register number that is not listed above might result in undefined behavior of the phy. This register is not applicable for clock lanes.

BIOS sets {D11ProcessFreqCtlOverride, D11ProcessFreqCtlIndex2} according to the following pseudocode at a Gen1 frequency and prior to the warm reset to transition to a Gen3 frequency:

- Temp[3:0] = Lookup in Table 238. The row is selected by the value of D18F0x[E8,C8,A8,88][Freq]. IF (D18F5x190[31:0]==0000_0000h) THEN the left value column of constants is selected ELSE the right column of values from D18F5x190 is selected.
- For each link, 0 to 3, and for each sublink if ungangled and the sublinks of the link differ in frequency:
 - IF (LinkGang || (Unganged and both sublinks at same frequency)) THEN
 - Sign = D18F4x1[9C,94,8C,84]_x4011_dm[1][FuseProcD11ProcessComp[2]]. (CADIN[0])
 - Offset[1:0] = D18F4x1[9C,94,8C,84]_x4011_dm[1][FuseProcD11ProcessComp[1:0]]. (CADIN[0])
 - ELSE // (unganged and both sublinks at different frequency)
 - Sign = D18F4x1[9C,94,8C,84]_x4[4,0]11_dm[1][FuseProcD11ProcessComp[2]]. (CADIN[8,0])
 - Offset[1:0] = D18F4x1[9C,94,8C,84]_x4[4,0]11_dm[1][FuseProcD11ProcessComp[1:0]]. (CADIN[8,0])
 - TempData[31:0] = D18F4x1[9C,94,8C,84]_x400F_dm[1][31:0]. (CADIN[0])
 - TempData[12] = 1. (D11ProcessFreqCtlOverride)
 - IF (Sign) THEN // Subtract offset
 - IF (Temp[3:0]<{00,Offset[1:0]}) THEN
 - AdjTemp[3:0] = 0h.
 - ELSE
 - AdjTemp[3:0] = (Temp[3:0]-{00,Offset[1:0]}).
 - ELSE // Sign==0, Add offset
 - IF ((Temp[3:0]+{00,Offset[1:0]})>Fh) THEN
 - AdjTemp[3:0] = Fh.
 - ELSE
 - AdjTemp[3:0] = (Temp[3:0]+{00,Offset[1:0]}).
 - TempData[3:0] = AdjTemp[3:0]. (D11ProcessFreqCtlIndex2)
 - IF (LinkGang || (Unganged and both sublinks at same frequency)) THEN

- D18F4x1[9C,94,8C,84]_x500F_dm[1][31:0] = TempData[31:0]. (Broadcast, 16-bit)
- ELSE // (unganged and both sublinks at different frequency)
 - D18F4x1[9C,94,8C,84]_x5[3:2]0F_dm[1][31:0] = TempData[31:0]. (Broadcast, 8-bit)

Table 238: BIOS Recommendation for DllProcessFreqCtlIndex2

Freq (GHz)	D18F5x190[31:0]	
	==0000_0000h	!=0000_0000h
1.2	Ah	D18F5x190[DllProcessFreqCtlIndex2Rate2X]
1.4		
1.6	7h	D18F5x190[DllProcessFreqCtlIndex2Rate3X]
1.8		
2.0	5h	D18F5x190[DllProcessFreqCtlIndex2Rate4X]
2.2		
2.4	4h	D18F5x190[DllProcessFreqCtlIndex2Rate48]
2.6	3h	D18F5x190[DllProcessFreqCtlIndex2Rate52]
2.8	3h	D18F5x190[DllProcessFreqCtlIndex2Rate56]
3.0	2h	D18F5x190[DllProcessFreqCtlIndex2Rate60]
3.2	2h	D18F5x190[DllProcessFreqCtlIndex2Rate64]

Bits	Description
31:13	Reserved.
12	DllProcessFreqCtlOverride . Read-write. Cold reset: 0. BIOS: See above for programming details. 1=Enables the override of DLL delay line capacitance settings with the values of DllProcessFreqCtlIndex2.
11:4	Reserved.
3:0	DllProcessFreqCtlIndex2 . Read-write. Cold reset: 0. BIOS: See above for programming details. Specifies the DLL delay line capacitance setting in lower power state.

D18F4x1[9C,94,8C,84]_x[5:4][9:0][9,1]1_dm[1] Link Phy Receiver Process Control

This is a direct map register set; see D18F4x1[98,90,88,80][DirectMapEn].

Table 239: Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][9,1]1_dm[1]

Register	Function
D18F4x1[9C,94,8C,84]_x4011_dm1	CADIN[0]
D18F4x1[9C,94,8C,84]_x4091_dm1	CADIN[1]
D18F4x1[9C,94,8C,84]_x4111_dm1	CADIN[2]
D18F4x1[9C,94,8C,84]_x4191_dm1	CADIN[3]
D18F4x1[9C,94,8C,84]_x4211_dm1	CADIN[4]
D18F4x1[9C,94,8C,84]_x4291_dm1	CADIN[5]
D18F4x1[9C,94,8C,84]_x4311_dm1	CADIN[6]
D18F4x1[9C,94,8C,84]_x4391_dm1	CADIN[7]

Table 239: Register Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][9,1]1_dm[1]

D18F4x1[9C,94,8C,84]_x4411_dm1	CADIN[8]
D18F4x1[9C,94,8C,84]_x4491_dm1	CADIN[9]
D18F4x1[9C,94,8C,84]_x4511_dm1	CADIN[10]
D18F4x1[9C,94,8C,84]_x4591_dm1	CADIN[11]
D18F4x1[9C,94,8C,84]_x4611_dm1	CADIN[12]
D18F4x1[9C,94,8C,84]_x4691_dm1	CADIN[13]
D18F4x1[9C,94,8C,84]_x4711_dm1	CADIN[14]
D18F4x1[9C,94,8C,84]_x4791_dm1	CADIN[15]
D18F4x1[9C,94,8C,84]_x4811_dm1	CTLIN[0]
D18F4x1[9C,94,8C,84]_x4911_dm1	CTLIN[1]

Table 240: Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[5:4][9:0][9,1]1_dm[1]

Link width	D18F4x1[98,90,88,80][15:8]	D18F4x1[98,90,88,80][7:0]
		11h
16-bit	50h	D18F4x1[8C,94,8C,84]_x[4[7:0][9,1]1,4[9:8]11]
8-bit	53h	D18F4x1[8C,94,8C,84]_x[4[7:4][9,1]1,4911]
8-bit	52h	D18F4x1[8C,94,8C,84]_x[4[3:0][9,1]1,4811]
4-bit	57h	D18F4x1[8C,94,8C,84]_x[4[5:4][9,1]1,4911]
4-bit	56h	D18F4x1[8C,94,8C,84]_x[4[1:0][9,1]1,4811]
2-bit	55h	D18F4x1[8C,94,8C,84]_x[44[9,1]1,4911]
2-bit	54h	D18F4x1[8C,94,8C,84]_x[40[9,1]1,4811]

Accessing any register number that is not listed above might result in undefined behavior of the phy. This register is not applicable for clock lanes.

Bits	Description
31:14	Reserved.
13:11	FuseProcDllProcessComp: fused DllProcessComp value. Read-write. Cold reset: Product-specific. See D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]F_dm[1][DllProcessFreqCtlIndex2].
10:0	Reserved.

D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]0_dm[1] Link Phy Transmit Control

This is a direct map register set; see D18F4x1[98,90,88,80][DirectMapEn]. The definition of this register is similar to that of D18F4x1[9C,94,8C,84]_x[D5,C5]. This register provides per-lane programmability whereas D18F4x1[9C,94,8C,84]_x[D5,C5] only allows per-sublink wide programmability.

Table 241: Register Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]0_dm[1]

Register	Function
D18F4x1[9C,94,8C,84]_x6000_dm1	CADOUT[0]
D18F4x1[9C,94,8C,84]_x6080_dm1	CADOUT[1]
D18F4x1[9C,94,8C,84]_x6100_dm1	CADOUT[2]

Table 241: Register Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]0_dm[1]

D18F4x1[9C,94,8C,84]_x6180_dm1	CADOUT[3]
D18F4x1[9C,94,8C,84]_x6200_dm1	CADOUT[4]
D18F4x1[9C,94,8C,84]_x6280_dm1	CADOUT[5]
D18F4x1[9C,94,8C,84]_x6300_dm1	CADOUT[6]
D18F4x1[9C,94,8C,84]_x6380_dm1	CADOUT[7]
D18F4x1[9C,94,8C,84]_x6400_dm1	CADOUT[8]
D18F4x1[9C,94,8C,84]_x6480_dm1	CADOUT[9]
D18F4x1[9C,94,8C,84]_x6500_dm1	CADOUT[10]
D18F4x1[9C,94,8C,84]_x6580_dm1	CADOUT[11]
D18F4x1[9C,94,8C,84]_x6600_dm1	CADOUT[12]
D18F4x1[9C,94,8C,84]_x6680_dm1	CADOUT[13]
D18F4x1[9C,94,8C,84]_x6700_dm1	CADOUT[14]
D18F4x1[9C,94,8C,84]_x6780_dm1	CADOUT[15]
D18F4x1[9C,94,8C,84]_x6800_dm1	CTLOUT[0]
D18F4x1[9C,94,8C,84]_x6880_dm1	CLKOUT[0]
D18F4x1[9C,94,8C,84]_x6900_dm1	CTLOUT[1]
D18F4x1[9C,94,8C,84]_x6980_dm1	CLKOUT[1]

Table 242: Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]0_dm[1]

Link width	D18F4x1[98,90,88,80][15:8]	D18F4x1[98,90,88,80][7:0]
		00h
16-bit	70h	D18F4x1[8C,94,8C,84]_x6[9:0][8,0]0
8-bit	73h	D18F4x1[8C,94,8C,84]_x6[9,7:4][8,0]0
8-bit	72h	D18F4x1[8C,94,8C,84]_x6[8,3:0][8,0]0
4-bit	77h	D18F4x1[8C,94,8C,84]_x6[9,5:4][8,0]0
4-bit	76h	D18F4x1[8C,94,8C,84]_x6[8,1:0][8,0]0
2-bit	75h	D18F4x1[8C,94,8C,84]_x6[9,4][8,0]0
2-bit	74h	D18F4x1[8C,94,8C,84]_x6[8,0][8,0]0

Accessing any register number that is not listed above might result in undefined behavior of the phy. These register bits are updated as specified by D18F0x16C[ImmUpdate].

Bits	Description
31:8	Reserved.
7	TxLs23ClkGateEn: LS2/LS3 clock gating enable. Read-write. Cold reset: 1. 1=Internal phy clock grids are gated during LS2 or PHY OFF states to save power. Recommended to set TxLs23ClkGateEn for systems that support LDTSTOP and LS2.
6:0	Reserved.

D18F4x1[9C,94,8C,84]_x6[9:8]84_dm[1] Link Phy Transmit Clock Phase Control

This is a direct map register set; see D18F4x1[98,90,88,80][DirectMapEn].

Table 243: Register Mapping for D18F4x1[9C,94,8C,84]_x6[9:8]84_dm[1]

Register	Function
D18F4x1[9C,94,8C,84]_x6884_dm1	CLKOUT[0]
D18F4x1[9C,94,8C,84]_x6984_dm1	CLKOUT[1]

Accessing any register number that is not listed above might result in undefined behavior of the phy. These register bits are updated as specified by D18F0x16C[ImmUpdate].

This register is not applicable for CAD/CTL lanes.

Bits	Description
31:1	Reserved.
0	Tx90En: transmit clock 90 degree phase shift enable. Read-write. Cold reset: 1. 1=Enable 90 degree phase shift on transmit clock lanes relative to CAD/CTL lanes. 0=Clock and CAD/CTL lanes are transmitted in phase. This feature might help reduce cross-talk.

D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]C_dm[1] Link Phy Tx Deemphasis and Margin Test Control

This is a direct map register set; see D18F4x1[98,90,88,80][DirectMapEn]. These registers specify the deemphasis values as well as provide test and debug features control associated with the link transmit lanes.

Table 244: Register Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]C_dm[1]

Register	Function
D18F4x1[9C,94,8C,84]_x600C_dm1	CADOUT[0]
D18F4x1[9C,94,8C,84]_x608C_dm1	CADOUT[1]
D18F4x1[9C,94,8C,84]_x610C_dm1	CADOUT[2]
D18F4x1[9C,94,8C,84]_x618C_dm1	CADOUT[3]
D18F4x1[9C,94,8C,84]_x620C_dm1	CADOUT[4]
D18F4x1[9C,94,8C,84]_x628C_dm1	CADOUT[5]
D18F4x1[9C,94,8C,84]_x630C_dm1	CADOUT[6]
D18F4x1[9C,94,8C,84]_x638C_dm1	CADOUT[7]
D18F4x1[9C,94,8C,84]_x640C_dm1	CADOUT[8]
D18F4x1[9C,94,8C,84]_x648C_dm1	CADOUT[9]
D18F4x1[9C,94,8C,84]_x650C_dm1	CADOUT[10]
D18F4x1[9C,94,8C,84]_x658C_dm1	CADOUT[11]
D18F4x1[9C,94,8C,84]_x660C_dm1	CADOUT[12]
D18F4x1[9C,94,8C,84]_x668C_dm1	CADOUT[13]
D18F4x1[9C,94,8C,84]_x670C_dm1	CADOUT[14]
D18F4x1[9C,94,8C,84]_x678C_dm1	CADOUT[15]
D18F4x1[9C,94,8C,84]_x680C_dm1	CTLOUT[0]
D18F4x1[9C,94,8C,84]_x688C_dm1	CLKOUT[0]
D18F4x1[9C,94,8C,84]_x690C_dm1	CTLOUT[1]
D18F4x1[9C,94,8C,84]_x698C_dm1	CLKOUT[1]

Table 245: Broadcast Mapping for D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]C_dm[1]

Link width	D18F4x1[98,90,88,80][15:8]	D18F4x1[98,90,88,80][7:0]
		0Ch
16-bit	70h	D18F4x1[8C,94,8C,84]_x6[9:0][8,0]C
8-bit	73h	D18F4x1[8C,94,8C,84]_x6[9,7:4][8,0]C
8-bit	72h	D18F4x1[8C,94,8C,84]_x6[8,3:0][8,0]C
4-bit	77h	D18F4x1[8C,94,8C,84]_x6[9,5:4][8,0]C
4-bit	76h	D18F4x1[8C,94,8C,84]_x6[8,1:0][8,0]C
2-bit	75h	D18F4x1[8C,94,8C,84]_x6[9,4][8,0]C
2-bit	74h	D18F4x1[8C,94,8C,84]_x6[8,0][8,0]C

Accessing any register number that is not listed above might result in undefined behavior of the phy.

This register specifies the deemphasis and voltage margining settings for the transmit drivers. See [2.12.3 \[Equalization\]](#) for more information about deemphasis. Links connecting internal nodes as specified by [D18F0x1A0\[IntLnkRoute\]](#) do not require deemphasis. These register bits are updated as specified by [D18F0x16C\[ImmUpdate\]](#).

For Gen3 links, the fields in this register can be programmed during link initialization to select the right deemphasis setting as follows.

Table 246: Gen3 Deemphasis Settings

Gen3 Deemphasis ¹ Setting	Post1	Pre1	Post2	Margin
No Deemphasis (0 dB)	0	0	0	0
-3 dB postcursor	38	0	0	0
-6 dB postcursor	64	0	0	0
-8 dB postcursor	77	0	0	0
-11 dB postcursor	92	0	0	0
-11 dB postcursor with -8 dB precursor ²	77	15	0	0

1. Deemphasis is not supported by the transmit clock lanes.
2. This setting is expected to be used for links with very high loss.

Deemphasis is not supported when operating at Gen1 link frequencies. Post1, Post2, and Pre1 must be 0 for Gen1 links.

[Table 247](#) and [Table 248](#) show recommended deemphasis settings for systems that have not been fully modeled and can also be used as estimates to begin margining work.

Table 247: Estimated deemphasis (dB) for microstrip routing

Length of trace (in)	Frequency of operation (GHz)					
	3.2	2.6	2.4	2.2	2.0	1.8
4	-3	0 ¹ /-3 ² /-3 ³	0 ¹ /0 ² /-3 ³	0	0	0
5	-3	-3	0 ¹ /-3 ² /-3 ³	0	0	0
6	-3	-3	-3	0	0	0

Table 247: Estimated deemphasis (dB) for microstrip routing

Length of trace (in)	Frequency of operation (GHz)					
	3.2	2.6	2.4	2.2	2.0	1.8
7	-3	-3	-3	$0^1/-3^2/-3^3$	0	0
8	-3	-3	-3	-3	0	0
9	-6	-3	-3	-3	0	0
10	-6	-3	-3	-3	0	0
11	-6	-3	-3	-3	-3	0
12	-6	-3	-3	-3	-3	-3
13	-6	$-3^1/-3^2/-6^3$	-3	-3	-3	-3
14	-6	$-3^1/-6^2/-6^3$	-3	-3	-3	-3
15	-6	-6	-3	-3	-3	-3
16	-6	-6	$-3^1/-6^2/-6^3$	-3	-3	-3
17	-6	-6	-6	-3	-3	-3
18	-8	-6	-6	-3	-3	-3
19	-8	-6	-6	$-3^1/-3^2/-6^3$	-3	-3
20	-8	-6	-6	-6	-3	-3

1. 0 connectors
2. 1 connector
3. 2 connectors

Table 248: Estimated deemphasis (dB) for stripline routing

Length of trace (in)	Frequency of operation (GHz)					
	3.2	2.6	2.4	2.2	2.0	1.8
4	-3	$0^1/-3^2/-3^3$	$0^1/0^2/-3^3$	0	0	0
5	-3	-3	-3	0	0	0
6	-3	-3	-3	$0^1/0^2/-3^3$	0	0
7	-3	-3	-3	-3	0	0
8	-3	-3	-3	-3	0	0
9	$-3^1/-6^2/-6^3$	-3	-3	-3	$0^1/0^2/-3^3$	0
10	$-3^1/-6^2/-6^3$	-3	-3	-3	-3	-3
11	-6	-3	-3	-3	-3	-3
12	-6	-3	-3	-3	-3	-3
13	-6	$-3^1/-3^2/-6^3$	-3	-3	-3	-3
14	-6	$-3^1/-3^2/-6^3$	-3	-3	-3	-3
15	-6	-6	$-3^1/-3^2/-6^3$	-3	-3	-3
16	-6	-6	-6	-3	-3	-3
17	-6	-6	-6	-3	-3	-3
18	$-6^1/-8^2/-8^3$	-6	-6	$-3^1/-3^2/-6^3$	-3	-3

Table 248: Estimated deemphasis (dB) for stripline routing

Length of trace (in)	Frequency of operation (GHz)					
	3.2	2.6	2.4	2.2	2.0	1.8
19	-8	-6	-6	-6	-3	-3
20	-8	-6	-6	-6	-3	-3

1. 0 connectors
2. 1 connector
3. 2 connectors

Bits	Description
31:24	Post2: deemphasis post-cursor level 2. Read-write. Cold reset: 0. Specifies the relative strength the post cursor 2 deemphasis tap. Value must be less than or equal to 104. BIOS: See Table 246 [Gen3 Deemphasis Settings] .
23:16	Post1: deemphasis post-cursor level 1. Read-write. Cold reset: 38. Specifies the relative strength the post cursor 1 deemphasis tap. Value must be less than or equal to 104. BIOS: See Table 246 [Gen3 Deemphasis Settings] .
15:8	Pre1: deemphasis pre-cursor level 1. Read-write. Cold reset: 0. Specifies the relative strength the pre cursor deemphasis tap. Value must be less than or equal to 104. BIOS: See Table 246 [Gen3 Deemphasis Settings] .
7:0	Margin. Read-write. Cold reset: 0. Specifies the relative strength the margin tap. Value must be less than or equal to 104. BIOS: See Table 246 [Gen3 Deemphasis Settings] .

D18F4x1B8 Processor TDP

See [MSRC001_0077](#) for converting TDP to watts.

Bits	Description
31:16	BaseTdp. Read-only. Value: Product-specific. Specifies the maximum TDP consumed by the processor for NB and logic external to the core. If D18F3xE8[MultiNodeCpu]=1 , this field reflects the maximum TDP consumed by the Northbridge and logic external to the core on both nodes. See MSRC001_0077[BaseTdp] .
15:0	ProcessorTdp. Read-only. Value: Product-specific. Specifies the maximum TDP the processor can support. See MSRC001_0077[ProcessorTdp] .

D18F4x1C4 L3 Power Control

Bits	Description
31:9	Reserved.

8	L3PwrSavEn: L3 power savings enable. Read-write. Reset: 0. BIOS: 1. 1=The clock to idle sub-caches in the L3 is stopped.																
7:0	<p>L3ClkHysCtl: L3 clock hysteresis control. Read-write. Reset: 10h. Specifies the maximum number of clocks that the clock is allowed to stop for an idle subcache. After the hysteresis counter expires the clock is restarted for one clock. If the subcache is still idle after sending one clock to the subcache, the clock is stopped and the hysteresis counter is restarted.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Reserved</td> </tr> <tr> <td>01h</td> <td>1</td> </tr> <tr> <td>02h</td> <td>2</td> </tr> <tr> <td>0Fh-03h</td> <td><L3ClkHysCtl></td> </tr> <tr> <td>10h</td> <td>16</td> </tr> <tr> <td>FEh-11h</td> <td>Reserved</td> </tr> <tr> <td>FFh</td> <td>Reserved</td> </tr> </tbody> </table> <p>BIOS should not change this field from the default state.</p>	<u>Bits</u>	<u>Description</u>	00h	Reserved	01h	1	02h	2	0Fh-03h	<L3ClkHysCtl>	10h	16	FEh-11h	Reserved	FFh	Reserved
<u>Bits</u>	<u>Description</u>																
00h	Reserved																
01h	1																
02h	2																
0Fh-03h	<L3ClkHysCtl>																
10h	16																
FEh-11h	Reserved																
FFh	Reserved																

D18F4x1C8 L3 Hit Statistics

Bits	Description												
31:28	Reserved.												
27:4	L3HitStat: L3 hit statistics. Read-only; Updated-by-hardware . Reset: X. See StatSel for the definition of the statistics returned by this field.												
3	StatSel: L3 hit statistics select. Read-write. Reset: 0. 1=L3HitStat[23:0] returns a moving average of the number of L3 hits per $6 * 2^{26}$ Northbridge clocks. 0=L3HitStat[9:0] returns a moving average of the number of L3 hits per 2^{10} L3 accesses.												
2:0	<p>ComputeUnitSel: compute unit select. Read-write. Reset: 0. Selects the compute unit for measuring L3 hit statistics.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Compute unit 0</td> </tr> <tr> <td>001b</td> <td>Compute unit 1</td> </tr> <tr> <td>010b</td> <td>Compute unit 2</td> </tr> <tr> <td>011b</td> <td>Compute unit 3</td> </tr> <tr> <td>111b-100b</td> <td>Reserved</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	000b	Compute unit 0	001b	Compute unit 1	010b	Compute unit 2	011b	Compute unit 3	111b-100b	Reserved
<u>Bits</u>	<u>Description</u>												
000b	Compute unit 0												
001b	Compute unit 1												
010b	Compute unit 2												
011b	Compute unit 3												
111b-100b	Reserved												

D18F4x1CC L3 Control 2

Read-write. Reset: 0000_0000h.

Bits	Description
31:9	Reserved.
8:6	ImplRdProjDelayThresh: implicit read projected delay threshold. BIOS: 010b. The enabled implicit selection paths are only delayed if the projected delay is less than the threshold.
5	Reserved.

4	ImplRdAnySubUnavail: implicit read if any subcache unavailable. BIOS: 1. 1=Apply the implicit read delay algorithm if any subcaches are unavailable. 0=Apply the implicit read delay algorithm only if all subcaches are unavailable.
3:0	Reserved.

D18F4x1D4 Compute Unit Based L3 Cache Partitioning

See 2.9.4.3 [L3 Cache Partitioning].

Bits	Description										
31:30	Reserved.										
29:26	MaskUpdateForComputeUnit: mask write to subcache enables for compute unit. Read-write. Reset: 0h. The value of MaskUpdateForComputeUnit for the current write directs which of ComputeUnit0SubCacheEn-ComputeUnit3SubCacheEn will be modified by the current write. 1=Mask write to the subcache enables for the specified logical compute unit. 0=Enable write to the subcache enables for the specified logical compute unit. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Mask write to ComputeUnit0SubCacheEn.</td> </tr> <tr> <td>[1]</td> <td>Mask write to ComputeUnit1SubCacheEn.</td> </tr> <tr> <td>[2]</td> <td>Mask write to ComputeUnit2SubCacheEn.</td> </tr> <tr> <td>[3]</td> <td>Mask write to ComputeUnit3SubCacheEn.</td> </tr> </tbody> </table>	Bit	Description	[0]	Mask write to ComputeUnit0SubCacheEn.	[1]	Mask write to ComputeUnit1SubCacheEn.	[2]	Mask write to ComputeUnit2SubCacheEn.	[3]	Mask write to ComputeUnit3SubCacheEn.
Bit	Description										
[0]	Mask write to ComputeUnit0SubCacheEn.										
[1]	Mask write to ComputeUnit1SubCacheEn.										
[2]	Mask write to ComputeUnit2SubCacheEn.										
[3]	Mask write to ComputeUnit3SubCacheEn.										
25:16	Reserved.										
15:12	ComputeUnit3SubCacheEn: compute unit 3 subcache enables. See: ComputeUnit0SubCacheEn. IF (MaskUpdateForComputeUnit[3]) THEN Read-only. ELSE Read-write. ENDIF. This field is reserved if there are < 4 compute units.										
11:8	ComputeUnit2SubCacheEn: compute unit 2 subcache enables. See: ComputeUnit0SubCacheEn. IF (MaskUpdateForComputeUnit[2]) THEN Read-only. ELSE Read-write. ENDIF. This field is reserved if there are < 3 compute units.										
7:4	ComputeUnit1SubCacheEn: compute unit 1 subcache enables. See: ComputeUnit0SubCacheEn. IF (MaskUpdateForComputeUnit[1]) THEN Read-only. ELSE Read-write. ENDIF. This field is reserved if there is 1 compute unit.										
3:0	ComputeUnit0SubCacheEn: compute unit 0 subcache enables. IF (MaskUpdateForComputeUnit[0]) THEN Read-only. ELSE Read-write. ENDIF. Reset: Fh. 1=Enable specified subcache for allocation of L2 victim. 0=Disable specified subcache for allocation for L2 victim; send to DRAM. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>Subcache 0</td> </tr> <tr> <td>[1]</td> <td>Subcache 1</td> </tr> <tr> <td>[2]</td> <td>Subcache 2</td> </tr> <tr> <td>[3]</td> <td>Subcache 3</td> </tr> </tbody> </table>	Bit	Description	[0]	Subcache 0	[1]	Subcache 1	[2]	Subcache 2	[3]	Subcache 3
Bit	Description										
[0]	Subcache 0										
[1]	Subcache 1										
[2]	Subcache 2										
[3]	Subcache 3										

3.8 Device [1F:18]h Function 5 Configuration Registers

See 3.1 [Register Descriptions and Mnemonics]. See 2.8 [Configuration Space].

D18F5x00 Device/Vendor ID

Bits	Description
31:16	DeviceID: device ID. Read-only. Value: 1605h.
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	See: MSRC001_024[6,4,2,0][31:0] .

D18F5x[74,64,54,44] Northbridge Performance Event Select High

Bits	Description
31:0	See: MSRC001_024[6,4,2,0][63:32] .

D18F5x[78,68,58,48] Northbridge Performance Event Counter Low

Bits	Description
31:0	See: MSRC001_024[7,5,3,1][31:0] .

D18F5x[7C,6C,5C,4C] Northbridge Performance Event Counter High

Bits	Description
31:0	See: MSRC001_024[7,5,3,1][63:32] .

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 249: 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.
7h	7h	3 compute units are enabled; both cores of each compute unit are enabled.
Fh	Fh	4 compute units are enabled; both cores of each compute unit are enabled.
1h	0h	1 compute unit is enabled; core 0 of the compute unit is enabled; core 1 of the compute unit is disabled.
3h	0h	2 compute units are enabled; core 0 of each compute unit is enabled; core 1 of each compute unit is disabled.
7h	0h	3 compute units are enabled; core 0 of each compute unit is enabled; core 1 of each compute unit is disabled.
Fh	0h	4 compute units are enabled; core 0 of each compute unit is enabled; core 1 of each compute unit is disabled.

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 249 [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>Compute unit 2</td> </tr> <tr> <td>[3]</td> <td>Compute unit 3</td> </tr> </tbody> </table>	Bit	Description	[0]	Compute unit 0	[1]	Compute unit 1	[2]	Compute unit 2	[3]	Compute unit 3
Bit	Description										
[0]	Compute unit 0										
[1]	Compute unit 1										
[2]	Compute unit 2										
[3]	Compute unit 3										
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 249 [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>Compute unit 2</td> </tr> <tr> <td>[3]</td> <td>Compute unit 3</td> </tr> </tbody> </table>	Bit	Description	[0]	Compute unit 0	[1]	Compute unit 1	[2]	Compute unit 2	[3]	Compute unit 3
Bit	Description										
[0]	Compute unit 0										
[1]	Compute unit 1										
[2]	Compute unit 2										
[3]	Compute unit 3										

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 165 [Valid Values for D18F2x94_dct[1:0][MemClkFreq]] ; 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 cores on the node is CmpCap+1-(the number of cores disabled by D18F3x190[DisCore]), as CmpCap is unaffected by D18F3x190[DisCore] . See 2.4.3 [Processor Cores and Downcoring] .

D18F5x88 NB Configuration 4 (NB_CFG4)

Bits	Description
31:10	Reserved.
9	DisHintInHtMskCnt. Read-write. Reset: 0. 1=Force HT Msk/Cnt to 0. 0=Allow HT Msk/Cnt to be non-0; used to contain core prefetch hints.

8:7	Reserved.
6	EnLnkTriStateNbPstate. Read-write. Cold Reset: 0. BIOS: 1. 1= D18F0x[E4,C4,A4,84][LdtStopTriEn] applies for LDTSTOPs for NB P-State changes. 0= D18F0x[E4,C4,A4,84][LdtStopTriEn] is ignored for LDTSTOPs for NB P-State changes; LDTSTOP_L assertion leaves the link transmitter signals in L0; receivers enabled after LDTSTOP_L deassertion based on internal timing. This field only applies to links operating at Gen1 frequencies.
5	Reserved.
4:0	Reserved.

D18F5xE0 Processor TDP Running Average

See [MSRC001_0078](#).

Bits	Description
31:26	Reserved.
25:4	TdpRunAvgAccCap: processor TDP running average accumulator capture. Read-only; updated-by-hardware. Reset: 0. Specifies the captured two's complement signed value of the processor TDP running average accumulator. The processor TDP running average accumulator is reset to 0 once the value is captured into TdpRunAvgAccCap or when D18F5xE8[ApmTdpLimit] changes. On an SCM, TdpRunAvgAccCap is a 21-bit two's complement signed value, sign extended to 22-bits. See RunAvgRange .
3:0	RunAvgRange: running average range. Read-write; Same-for-all . Reset: 0. BIOS: 9h. Specifies the bit that the internal FreeRunSampleTimer must overflow to cause the processor TDP running average accumulator value to be captured into TdpRunAvgAccCap. The time interval is specified as $2^{(\text{RunAvgRange} + 1)} * \text{FreeRunSampleTimer rate}$. A value of 0 disables the TDP running average accumulator capture function. See 2.5.2.1.1 [Application Power Management (APM)] .

D18F5xE8 TDP Limit 3

Bits	Description
31:29	Reserved.
28:16	ApmTdpLimit. Read-only; updated-by-hardware. If the consumed node power exceeds the ApmTdpLimit on a single node processor or the ApmTdpLimit/2 on a multi-node processor, a P-state limit is applied to all cores on all nodes to reduce the power consumption to remain within the TDP limit. See 2.5.2.1.1.1 [TDP Limiting] and MSRC001_0078[ApmTdpLimit] . Value: $(\text{MSRC001_0075}[\text{ApmTdpLimit}] == 0)$ $? (\text{D18F4x10C}[\text{NodeTdpLimit}] * (\text{D18F3xE8}[\text{MultiNodeCpu}] + 1))$ $: \text{MIN}((\text{D18F4x10C}[\text{NodeTdpLimit}] * (\text{D18F3xE8}[\text{MultiNodeCpu}] + 1)),$ $\text{MSRC001_0075}[\text{ApmTdpLimit}]).$
15:10	Tdp2Watt[5:0]. Read-only. Value: 000000b. See Tdp2Watt[15:6] .
9:0	Tdp2Watt[15:6]. Read-only. Value: Product-specific. Specifies in watts/TDP units the conversion factor for converting TDP units to watts. Tdp2Watt[15:0] is a fixed point integer with 16 bits to the right of the decimal point and 0 bits to the left of the decimal point. E.g. Tdp2Watt[15] ==0.5 W; Tdp2Watt[6] ==0.976 mW; Tdp2Watt[0] ==15.2 uW.

D18F5x128 Clock Power/Timing Control 3

Bits	Description												
31:17	Reserved.												
16	CC6PwrDwnVcoEn: CC6 power down VCO enable. Read-write. Cold reset: Product-specific. 1=Power down the VCO on CC6 entry. If this bit is set to 1, then CC6PwrDwnRegEn must be 0.												
15	CC6PwrDwnRegEn: CC6 power down regulator enable. Read-write. Cold reset: Product-specific. 1=Power down the VDDA regulator on CC6 entry. If this bit is set to 1, then CC6PwrDwnVcoEn must be 0. See PllRegTime.												
14	Reserved.												
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. See 2.12.7 [LDTSTOP_L Requirements].</td> </tr> <tr> <td>01b</td> <td>1 us</td> <td>11b</td> <td>10 us. See 2.12.7 [LDTSTOP_L Requirements].</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	00b	500 ns	10b	5 us. See 2.12.7 [LDTSTOP_L Requirements].	01b	1 us	11b	10 us. See 2.12.7 [LDTSTOP_L Requirements].
Bits	Description	Bits	Description										
00b	500 ns	10b	5 us. See 2.12.7 [LDTSTOP_L Requirements].										
01b	1 us	11b	10 us. See 2.12.7 [LDTSTOP_L Requirements].										
11:9	Reserved.												
8:7	PllRegTime: PLL regulator time. Read-write. Cold Reset: 10b. BIOS: 10b. The VDDAregulator may be powered down when the processor transitions to CC6. See CC6PwrDwnRegEn. 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. <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>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>	Bits	Description	Bits	Description	00b	Reserved.	10b	1.5 us	01b	Reserved.	11b	2.0 us
Bits	Description	Bits	Description										
00b	Reserved.	10b	1.5 us										
01b	Reserved.	11b	2.0 us										
6:0	Reserved.												

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 250: Register Mapping for D18F5x1[6C:60]

Register	Function
D18F5x160	NB P-state 0
D18F5x164	NB P-state 1
D18F5x168	Reserved
D18F5x16C	Reserved

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 2.5.2.2 [NB P-states] for more information about these registers.

Table 251: NB P-state Definitions

Term	Definition
NBCOF	NB current operating frequency in MHz. $NBCOF = 200 * (D18F5x1[6C:60][NbFid] + 4h) / (2^{D18F5x1[6C:60][NbDid]})$.
NBCOF[0]	NB current operating frequency in MHz for NB P-state 0. $NBCOF[0] = (200 * (D18F5x160[NbFid] + 4h) / (2^{D18F5x160[NbDid]}))$.
NBCOF[1]	NB current operating frequency in MHz for NB P-state 1. $NBCOF[1] = (200 * (D18F5x164[NbFid] + 4h) / (2^{D18F5x164[NbDid]}))$.

Bits	Description
31:17	Reserved.
16:10	NbVid: NB VID. Read-write. Specifies the NB voltage. Writes outside the MSRC001_0071 [MaxVid, MinVid] range are ignored.
9:8	Reserved.
7	NbDid: NB divisor ID. Read-write; Same-for-all. Specifies the NB frequency divisor; see NbFid.
6	Reserved.
5:1	NbFid: NB frequency ID. Read-write; Same-for-all. Specifies the NB 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 [MaxNbCof]. NbFid must be less than or equal to 1Bh.
0	NbPstateEn: NB P-state enable. Read-write; Same-for-all. 1=The NB P-state specified by this register is valid. 0=The NB P-state specified by this register is not valid. See D18F5x170 [NbPstateMaxVal].

D18F5x170 Northbridge P-state Control

See also [2.5.2.2 \[NB P-states\]](#).

Bits	Description
31:16	Reserved.
15	Reserved.
14	SwNbPstateLoDis: software NB P-state low disable. IF (MSRC001_0071 [NbPstateDis]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0. BIOS: 0. 1=Transition to NbPstateHi and disable transitions to NbPstateLo.
13	NbPstateDisOnP0: NB P-state disable on P0. IF (MSRC001_0071 [NbPstateDis]) THEN Read-only. ELSE Read-write. ENDIF. Reset: 0. BIOS: 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 2.5.2.1.2.1 [Software P-state Numbering] .
12	Reserved.
11:9	NbPstateThreshold: NB P-state threshold. Read-write. Reset: Product-specific. BIOS: NumOf-CompUnitsOnNode . 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.
8	Reserved.

7:6	NbPstateHi: NB P-state high. Read-write; Same-for-all . 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 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] .
5	Reserved.
4:3	NbPstateLo: NB P-state low. Read-write; Same-for-all . 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 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] .
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										
31:21	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" data-bbox="267 1323 665 1501"> <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: current northbridge voltage ID. See: MSRC001_0071[CurNbVid] .										
11:10	Reserved.										
9	CurNbDid: current northbridge divisor ID. Read-only. Cold reset: Product-specific.										
8	Reserved.										
7:3	CurNbFid: current northbridge frequency ID. Read-only. Cold reset: Product-specific.										
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. See MSRC001_0071[NbPstateDis] .										

D18F5x190 Link Product Information

Bits	Description
31:28	DllProcessFreqCtlIndex2Rate64 . Value: Product-specific.
27:24	DllProcessFreqCtlIndex2Rate60 . Value: Product-specific.
23:20	DllProcessFreqCtlIndex2Rate56 . Value: Product-specific.
19:16	DllProcessFreqCtlIndex2Rate52 . Value: Product-specific.
15:12	DllProcessFreqCtlIndex2Rate48 . Value: Product-specific.
11:8	DllProcessFreqCtlIndex2Rate4X . Value: Product-specific.
7:4	DllProcessFreqCtlIndex2Rate3X . Value: Product-specific.
3:0	DllProcessFreqCtlIndex2Rate2X . Value: Product-specific.

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" data-bbox="267 1087 1096 1188"> <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. Provides the recommended settings for MSRC001_00[35:30] [Processor Name String].

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 .

3.9 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	<p>ApicId: APIC ID. Read-write. Reset: Varies based on core number and node number.</p> <ul style="list-style-type: none"> • If \simD18F3xE8[MultiNodeCpu] then the initial value of ApicId[7:0]={0b, D18F0x60[NodeId[2:0]], CpuCoreNum[3:0]}. • If D18F3xE8[MultiNodeCpu] then the initial value of ApicId[7:0]={0b, D18F0x60[NodeId[2:1]], CpuCoreNum[4:0]}. <p>See 2.4.8.1.3 [ApicId Enumeration Requirements]. See 2.4.3 [Processor Cores and Downcoring].</p>
23:0	Reserved.

APIC30 APIC Version

Bits	Description
31	ExtApicSpace: extended APIC register space present. Read-only. Reset: 1. 1=Indicates the presence of extended APIC register space starting at APIC400.
30:25	RAZ.
24	DirectedEoiSupport: directed EOI support. Read-only. Reset: 0. 0=Directed EOI capability not supported.
23:16	MaxLvtEntry. Read-only. Reset: Product-specific. Specifies the number of entries in the local vector table minus one.
15:8	RAZ.
7:0	Version. Read-only. 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. Reset: 0. Read-only. 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

Bits	Description
31:0	RemoteReadData. Read-only. Reset: 0. The data resulting from a valid completion of a remote read inter-processor interrupt.

APICD0 Logical Destination (LDR)

Bits	Description
31:24	Destination. Read-write. Reset: 0. This APIC's destination identification. Used to determine which interrupts should be accepted.
23:0	Reserved.

APICE0 Destination Format

Only supported in xAPIC mode.

Bits	Description								
31:28	Format. Read-write. Reset: Fh. Controls which format to use when accepting interrupts with a logical destination mode. <table border="1" data-bbox="267 1654 933 1795"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Logical Cluster Destination Mode.</td> </tr> <tr> <td>Eh-1h</td> <td>Reserved</td> </tr> <tr> <td>Fh</td> <td>Logical Flat Destination Mode.</td> </tr> </tbody> </table>	Bits	Description	0h	Logical Cluster Destination Mode.	Eh-1h	Reserved	Fh	Logical Flat Destination Mode.
Bits	Description								
0h	Logical Cluster Destination Mode.								
Eh-1h	Reserved								
Fh	Logical Flat Destination Mode.								
27:0	Reserved. Reset: FFF_FFFFh.								

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, and Startup interrupts may be accepted; pending interrupts in APIC[170:100] and APIC[270:200] are held, but further fixed, lowest-priority, LINT[1:0], 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)

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 252: [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. Reset: 0. Read-only. These bits are set when the corresponding interrupt is being serviced by the core.

APIC[1F0:180] Trigger Mode (TMR)

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 253: 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. Reset: 0. Read-only. 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)

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 254: 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. Read-only. Reset: 0. 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 255: 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. Reserved in x2APIC mode. 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_0071\[CurPstateLimit\]](#) due to SB-RMI, software P-state limit, 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 256: 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	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.

APIC380 Timer Initial Count

Bits	Description
31:0	Count. Read-write. Reset: 0. The value copied into the current count register when the timer is loaded or reloaded.

APIC390 Timer Current Count

Bits	Description
31:0	Count. Read-only. Reset: 0. The current value of the counter.

APIC3E0 Timer Divide Configuration

The Div bits are encoded as follows:

Table 257: 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 257 .
2	RAZ.
1:0	Div[1:0] . Read-write. Reset: 0. See Table 257 .

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 IntDest[7:0]=1111_1111b (instead of xxxx_1111b); a match in physical destination mode occurs when (IntDest[7:0] == ApicId[7:0]) instead of (IntDest[3:0] == ApicId[3:0]) . If 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

Interrupt enables range is mapped as follows:

Table 258: [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. Reset: FFFF_FFFFh. 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](#). APIC510 provides a local vector table entry for thresholding; See [D18F3x160](#), [D18F3x168](#), and [D18F3x170](#). APIC510 provides a local vector table entry for thresholding; See [D18F3x160](#), [D18F3x168](#), and [D18F3x170](#). APIC520 is unused. APIC530 provides a local vector table entry for SBI; See [D18F3x1E4](#).

Table 259: [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	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.

3.10 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 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 260: [Reset Mapping](#) for [CPUID Fn0000_0000_E\[D,C,B\]X](#)

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]=06h, then Family[7:0]=15h.

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 as CPUID Fn8000_0008_ECX[NC] + 1. Otherwise, this field is reserved. Value: Product-specific.
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: IF (PROC>=OR_C0) THEN 1 ELSE 0 ENDIF0.
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.

Bits	Description
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: IF (PROC>=OR_C0) THEN 1 ELSE 0 ENDIF0.
11:10	Reserved.
9	SSSE3: supplemental SSE3 extensions. Value: 1.
8:4	Reserved.
3	Monitor: Monitor/Mwait instructions. Value: ~MSRC001_0015[MonMwaitDis].
2	Reserved.
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].

Bits	Description
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)

Bits	Description
31:4	Reserved.
3	BMI1: bit manipulation group 1 instruction support. Value: IF (PROC >= OR_C0) THEN 1 ELSE 0 ENDIF0.
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[C:8] 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 261: Reset Mapping for CPUID Fn8000_0000_E[D,C,B]X

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

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]=06h, then Family[7:0]=15h.

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 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>Reserved</td> </tr> <tr> <td>0001b</td> <td>AM3r2</td> </tr> <tr> <td>0010b</td> <td>Reserved</td> </tr> <tr> <td>0011b</td> <td>G34r1</td> </tr> <tr> <td>0100b</td> <td>Reserved</td> </tr> <tr> <td>0101b</td> <td>C32r1</td> </tr> <tr> <td>1111b-0110b</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0000b	Reserved	0001b	AM3r2	0010b	Reserved	0011b	G34r1	0100b	Reserved	0101b	C32r1	1111b-0110b	Reserved
Bits	Description																
0000b	Reserved																
0001b	AM3r2																
0010b	Reserved																
0011b	G34r1																
0100b	Reserved																
0101b	C32r1																
1111b-0110b	Reserved																
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: 1. Indicates support for CPUID Fn8000_001D_EAX_x0 - CPUID Fn8000_001E_EDX .
21	TBM: trailing bit manipulation instruction support. Value: IF (PROC>=OR_C0) THEN 1 ELSE 0 ENDIF0.
20	Reserved.
19	Reserved.
18	Reserved.
17	TCE: translation cache extension. Value: IF (PROC>=OR_C0) THELSE 0 ENDIF0.
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.
3	ExtApicSpace: extended APIC register space. Value: 1.
2	SVM: Secure Virtual Mode feature. Value: Product-specific. Indicates support for: VMRUN, VMLOAD, VMSAVE, CLGI, VMPCALL, 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.
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 262: Reset Mapping for CPUID Fn8000_000[4:2]_E[D,C,B,A]X

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_ECX[L1DcAssoc].
23:16	L1DTlb2and4MSize: data TLB number of entries for 2 MB and 4 MB pages. Value: IF (PROC>=OR_C0) THEN 64 ELSE 32 ENDIF32. 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_ECX[L1DcAssoc].
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_ECX [L1DcAssoc].
23:16	L1DTlb4KSize: data TLB number of entries for 4 KB pages. Value: IF (PROC>=OR_C0) THEN 64 ELSE 32 ENDIF32.
15:8	L1ITlb4KAssoc: instruction TLB associativity for 4 KB pages. Value: FFh. See: CPUID Fn8000_0005_ECX [L1DcAssoc].
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 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. Value: 4. <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>1 way (direct mapped)</td> </tr> <tr> <td>02h</td> <td>2 way</td> </tr> <tr> <td>03h</td> <td>3 way</td> </tr> <tr> <td>FEh-04h</td> <td><L1IcAssoc> way</td> </tr> <tr> <td>FFh</td> <td>Fully associative</td> </tr> </tbody> </table>	Bits	Description	00h	Reserved	01h	1 way (direct mapped)	02h	2 way	03h	3 way	FEh-04h	<L1IcAssoc> way	FFh	Fully associative
Bits	Description														
00h	Reserved														
01h	1 way (direct mapped)														
02h	2 way														
03h	3 way														
FEh-04h	<L1IcAssoc> way														
FFh	Fully associative														
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 cache characteristics shared by both cores of a compute unit.

Bits	Description
31:24	L1IcSize: L1 instruction cache size KB. Value: 64.
23:16	L1IcAssoc: L1 instruction cache associativity. Value: 2. See: CPUID Fn8000_0005_ECX [L1DcAssoc].
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 4K pages shared by each core on a compute 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 both cores of a compute unit.

Bits	Description																																				
31:16	L2Size: L2 cache size in KB. Value: Product-specific. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</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>	Bits	Description	03FFh-0000h	Reserved	0400h	1 MB	07FFh-0401h	Reserved	0800h	2 MB	FFFFh-0801h	Reserved																								
Bits	Description																																				
03FFh-0000h	Reserved																																				
0400h	1 MB																																				
07FFh-0401h	Reserved																																				
0800h	2 MB																																				
FFFFh-0801h	Reserved																																				
15:12	L2Assoc: L2 cache associativity. Value: 8. <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>Disabled.</td> <td>8h</td> <td>16 ways</td> </tr> <tr> <td>1h</td> <td>1 way (direct mapped)</td> <td>9h</td> <td>Reserved</td> </tr> <tr> <td>2h</td> <td>2 ways</td> <td>Ah</td> <td>32 ways</td> </tr> <tr> <td>3h</td> <td>Reserved</td> <td>Bh</td> <td>48 ways</td> </tr> <tr> <td>4h</td> <td>4 ways</td> <td>Ch</td> <td>64 ways</td> </tr> <tr> <td>5h</td> <td>Reserved</td> <td>Dh</td> <td>96 ways</td> </tr> <tr> <td>6h</td> <td>8 ways</td> <td>Eh</td> <td>128 ways</td> </tr> <tr> <td>7h</td> <td>Reserved</td> <td>Fh</td> <td>Fully associative</td> </tr> </tbody> </table>	Bits	Description	Bits	Description	0h	Disabled.	8h	16 ways	1h	1 way (direct mapped)	9h	Reserved	2h	2 ways	Ah	32 ways	3h	Reserved	Bh	48 ways	4h	4 ways	Ch	64 ways	5h	Reserved	Dh	96 ways	6h	8 ways	Eh	128 ways	7h	Reserved	Fh	Fully associative
Bits	Description	Bits	Description																																		
0h	Disabled.	8h	16 ways																																		
1h	1 way (direct mapped)	9h	Reserved																																		
2h	2 ways	Ah	32 ways																																		
3h	Reserved	Bh	48 ways																																		
4h	4 ways	Ch	64 ways																																		
5h	Reserved	Dh	96 ways																																		
6h	8 ways	Eh	128 ways																																		
7h	Reserved	Fh	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 function provides third level cache characteristics shared by all cores.

Bits	Description																										
31:18	<p>L3Size: L3 cache size. Value: Product-specific. The L3 cache size in 512 KB units. This value is reduced by the amount consumed by the probe filter; See 2.9.4.1 [Probe Filter]. This value reflects the total L3 cache for both nodes of a G34r1 processor. L3 cache size for the processor is at least (L3Size[31:18] * 512 KB) and less than ((L3Size[31:18] + 1) * 512 KB).</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0000h</td> <td>Disabled.</td> </tr> <tr> <td>0001h</td> <td>0.5 MB</td> </tr> <tr> <td>0002h</td> <td>1 MB</td> </tr> <tr> <td>0003h</td> <td><L3Size*0.5> MB</td> </tr> <tr> <td>0004h</td> <td>2 MB</td> </tr> <tr> <td>0007h-0005h</td> <td><L3Size*0.5> MB</td> </tr> <tr> <td>0008h</td> <td>4 MB</td> </tr> <tr> <td>000Fh-0009h</td> <td><L3Size*0.5> MB</td> </tr> <tr> <td>0010h</td> <td>8 MB</td> </tr> <tr> <td>001Fh-0011h</td> <td><L3Size*0.5> MB</td> </tr> <tr> <td>0020h</td> <td>16 MB</td> </tr> <tr> <td>3FFFh-0021h</td> <td>Reserved.</td> </tr> </tbody> </table>	Bits	Description	0000h	Disabled.	0001h	0.5 MB	0002h	1 MB	0003h	<L3Size*0.5> MB	0004h	2 MB	0007h-0005h	<L3Size*0.5> MB	0008h	4 MB	000Fh-0009h	<L3Size*0.5> MB	0010h	8 MB	001Fh-0011h	<L3Size*0.5> MB	0020h	16 MB	3FFFh-0021h	Reserved.
Bits	Description																										
0000h	Disabled.																										
0001h	0.5 MB																										
0002h	1 MB																										
0003h	<L3Size*0.5> MB																										
0004h	2 MB																										
0007h-0005h	<L3Size*0.5> MB																										
0008h	4 MB																										
000Fh-0009h	<L3Size*0.5> MB																										
0010h	8 MB																										
001Fh-0011h	<L3Size*0.5> MB																										
0020h	16 MB																										
3FFFh-0021h	Reserved.																										
17:16	Reserved.																										
15:12	<p>L3Assoc: L3 cache associativity. Value: Product-specific. This value reflects the total L3 associativity for the processor. For G34r1 processors, this value reflects the combined associativity of both nodes. The actual number of ways available to the program may be smaller than the value reported in this field, for example, when probe filter is enabled. Software should use CPUID CPUID Fn8000_001D_EBX_x3[CacheNumWays] for an accurate value. See CPUID Fn8000_0006_ECX[L2Assoc].</p>																										
11:8	L3LinesPerTag: L3 cache lines per tag. Value: 1.																										
7:0	L3LineSize: L3 cache line size in bytes. Value: 64.																										

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	Reserved.

10	EffFreqRO: read-only effective frequency interface. Value: IF (PROC>=OR_C0) THEN 1 ELSE 0 ENDIF. 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 2.5.2.1.1 [Application Power Management (APM)].
8	TscInvariant: TSC invariant. Value: 1. The TSC rate is invariant.
7	Hwpstate: 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: IF (D18F3xE8[MultiNodeCpu]) THEN 5h ELSE 4h ENDIF. The number of bits in the initial APIC20[ApicId] value that indicate core ID within a processor.

Bits	Description
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. This value is affected by D18F3x190[DisCore] . 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.
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 first level TLB characteristics for 1G pages shared by each core on a compute unit.

Bits	Description
31:28	L1DTlb1GAssoc: L1 data TLB associativity for 1 GB pages. Value: Fh. See: CPUID Fn8000_0006_ECX[L2Assoc] .
27:16	L1DTlb1GSize: L1 data TLB number of entries for 1 GB pages. Value: IF (PROC>=OR_C0) THEN 64 ELSE 32 ENDIF32.
15:12	L1ITlb1GAssoc: L1 instruction TLB associativity for 1 GB pages. Value: Fh. See: CPUID Fn8000_0006_ECX[L2Assoc] .
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 second level TLB characteristics for 1G pages shared by each core on a compute unit.

Bits	Description
31:28	L2DTlb1GAssoc: L2 data TLB associativity for 1 GB pages. Value: 6. See: CPUID Fn8000_0006_ECX[L2Assoc] .
27:16	L2DTlb1GSize: L2 data TLB number of entries for 1 GB pages. Value: 1024.

Bits	Description
15:12	L2ITlb1GAssoc: L2 instruction TLB associativity for 1 GB pages. Value: 6. See: CPUID Fn8000_0006_ECX[L2Assoc] .
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.
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: 80h. Offset from the start of the LWPCB to the EventInterval1 field.
23:16	LwpMaxEvents: maximum EventId. Value: 3. 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 0See [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.
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 0See [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 263: 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	L3

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.												
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. IF (PROC>=OR_C0) THEN Value: IF (D18F5x80[DualCore]==1) THEN 001h ELSE 000h ENDIF. See: CPUID Fn8000_001D_EAX_x0[NumSharingCache]. ELSE Value: 001h. See: CPUID Fn8000_001D_EAX_x0[NumSharingCache]. ENDIF.
13:10	Reserved.

Bits	Description
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. IF (PROC>=OR_C0) THEN Value: IF (D18F5x80[DualCore]==1) THEN 001h ELSE 000h ENDIF. See: CPUID Fn8000_001D_EAX_x0[NumSharingCache] . ELSE Value: 001h. See: CPUID Fn8000_001D_EAX_x0[NumSharingCache] . ENDIF.
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

[CPUID Fn8000_001D_EAX_x3](#) reports topology information for the L3. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:26	Reserved.
25:14	NumSharingCache: number of cores sharing cache. Value: D18F5x84[CmpCap]-NumOnes(D18F3x190[DisCore]). 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] .

Bits	Description
7:5	CacheLevel: cache level. Identifies the cache level. Value: 011b. 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_x4 Cache Properties

[CPUID Fn8000_001D_EAX_x4](#) reports done/null. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:5	Reserved. Value: 0.
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 [Cache-LineSize].

CPUID Fn8000_001D_EBX_x3 Cache Properties

[CPUID Fn8000_001D_EBX_x3](#) reports topology information for the L3. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:22	CacheNumWays: cache number of ways. Read-only; Same-within-pkg. Value: Product-specific. See: CPUID Fn8000_001D_EBX_x0 [CacheNumWays]. <ul style="list-style-type: none"> Note: This value reflects the number of ways available to the processor core, unlike CPUID Fn8000_0006_EDX[L3Assoc] which reports the total number of ways available in the complete L3 cache. The reported CacheNumWays accounts for cache ways that are not available to the processor core due to the split nature of the L3 on G34r1 processors, probe filter and/or compute unit based L3 cache partitioning. See 2.9.4.1 [Probe Filter] and 2.9.4.3 [L3 Cache Partitioning]. This value is not necessarily the same for all processor cores within a package if L3 cache partitioning is not evenly distributed amongst all compute units.
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 [Cache-LineSize].

CPUID Fn8000_001D_EBX_x4 Cache Properties

[CPUID Fn8000_001D_EBX_x4](#) reports done/null. See [CPUID Fn8000_001D_EAX_x0](#).

Bits	Description
31:0	Reserved. Value: 0000_0000h.

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 CacheNum-Sets+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

CPUID Fn8000_001D_ECX_x3 reports topology information for the L3. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:0	CacheNumSets: cache number of sets. Value: 0000_07FFh. See: CPUID Fn8000_001D_ECX_x0[CacheNumSets].

CPUID Fn8000_001D_ECX_x4 Cache Properties

CPUID Fn8000_001D_ECX_x4 reports done/null. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:0	Reserved. Value: 0000_0000h.

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

CPUID Fn8000_001D_EDX_x3 reports topology information for the L3. 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_x4 Cache Properties

CPUID Fn8000_001D_EDX_x4 reports done/null. See CPUID Fn8000_001D_EAX_x0.

Bits	Description
31:0	Reserved. Value: 0000_0000h.

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 IdentifiersSee [CPUID Fn8000_001E_EAX](#).

Bits	Description
31:16	Reserved.
15:8	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 IdentifiersSee [CPUID Fn8000_001E_EAX](#).

Bits	Description								
31:11	Reserved.								
10:8	NodesPerProcessor. Value: {00b,D18F3xE8[MultiNodeCpu]}. Specifies the number of nodes per processor. <table border="1" data-bbox="267 892 779 1039"> <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								
7:0	NodeId. Value: {00000b,D18F0x60[NodeId]}. Specifies the node ID.								

CPUID Fn8000_001E_EDX ReservedSee [CPUID Fn8000_001E_EAX](#).

Bits	Description
31:0	Reserved.

3.11 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

Unlike [MSR0000_0401](#), the access type of fields of [MSR0000_0001](#) are not affected by [MSRC001_0015](#)[McStatusWrEn]; [MSR0000_0401](#) is always writable.

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 2.5.2.1.2.1 [Software P-state Numbering] . The TSC counts at the same rate in all P-states, all C states, S0, or S1. A read of this MSR in guest mode 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. ApicEn[1:0] = {Apicx2En, ApicEn}. See 2.4.8.1.2 [APIC Register Space] . <table border="1" data-bbox="256 1171 990 1344"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Local APIC is disabled.</td> </tr> <tr> <td>01b</td> <td>Local APIC is enabled in xAPIC mode.</td> </tr> <tr> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>Local APIC is enabled in x2APIC mode.</td> </tr> </tbody> </table>	Bits	Description	00b	Local APIC is disabled.	01b	Local APIC is enabled in xAPIC mode.	10b	Reserved	11b	Local APIC is enabled in x2APIC mode.
Bits	Description										
00b	Local APIC is disabled.										
01b	Local APIC is enabled in xAPIC mode.										
10b	Reserved										
11b	Local APIC is enabled in x2APIC mode.										
10	Apicx2En: APIC enable. MBZ. Reset: 0. See ApicEn.										
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)

Read; GP-write. Writes to this register result in a GP faults with error code 0.

Bits	Description
63:18	MBZ.
17:16	ClusterID. Reset: 00b. The 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 register is used by software to determine the effective frequency of the core. A read of this MSR in guest mode is affected by MSRC000_0104 [Time Stamp Counter Ratio (TscRateMsr)] . This field uses software P-state numbering. See MSRC001_0015[EffFreqCntMwait] , 2.5.6 [Effective Frequency Interface] , and 2.5.2.1.2.1 [Software P-state Numbering] .

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 .

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: 0000h. 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 2.13.1 [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 2.13.1 [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 2.13.1 [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)

Read; GP-write; *Not-same-for-all*; *Updated-by-hardware*. Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	LastBranchFromIP. Loaded with the segment offset of the branch instruction.

MSR0000_01DC Last Branch To IP (BR_TO)

Read; GP-write; *Not-same-for-all*; *Updated-by-hardware*. Reset: 0000_0000_0000_0000h.

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

Read; GP-write; *Not-same-for-all*; *Updated-by-hardware*. Reset: 0000_0000_0000_0000h.

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

Read; GP-write; **Not-same-for-all**; **Updated-by-hardware**. Reset: 0000_0000_0000_0000h.

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.). Variable 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 264: [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 265: [Valid Values](#) for Memory Type

Bits	Description
000b	UC or uncacheable.
001b	WC or write combining.
011b-010b	Reserved
100b	WT or write through.

Table 265: [Valid Values](#) for Memory Type

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 265 [Valid Values for Memory Type] .

MSR0000_020[F,D,B,9,7,5,3,1] Variable-Size MTRRs Mask

SharedC.

Table 266: [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. 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 267: 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 268: 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. IF (MSRC001_0010 [MtrrFixDramModEn]) THEN Read-write. ELSE MBZ. ENDIF. Reset: 0. 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. IF (MSRC001_0010 [MtrrFixDramModEn]) THEN Read-write. ELSE MBZ. ENDIF. Reset: 0. 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. Reset: 0. Address range from 00000h to 0FFFFh. See: Table 265 [Valid Values for Memory Type] .

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. See [MSR0000_020\[F:0\]](#) for general MTRR information.

Bits	Description
63:12	MBZ.
11	MtrrDefTypeEn: variable and fixed MTRR enable . Read-write. Reset: 0. 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. Reset: 0. 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. Reset: 0. 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 2.13.1 [[Machine Check Architecture](#)]. See [MSRC001_0044](#) [[LS Machine Check Control Mask \(MC0_CTL_MASK\)](#)].

Bits	Description
63:10	Unused.
9	IntErrType1: internal error type 1 .
8	IntErrType2: internal error type 2 .
7	SRDE: read data errors . System read data errors on cache fill.

6	LineFillPoison: line fill poison error.
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 error.

MSR0000_0401 LS Machine Check Status (MC0_STATUS)

Cold reset: 0. See 2.13.1 [Machine Check Architecture]. See MSRC001_0015[McStatusWrEn]. Table 269 describes each error type. Table 270 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. Deferred errors 3. Correctable errors <p>The machine check mechanism handles the contents of MC_i_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, as indicated by MC _i _CTL.
59	MiscV: miscellaneous error register valid. Read-write; Updated-by-hardware. 1=Valid thresholding in MSR0000_0403.
58	AddrV: error address valid. Read-write; Updated-by-hardware. 1=MC _i _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 2.13.1.6.1 [MCA Differentiation Between System-Fatal and Process-Fatal Errors].

56:45	Reserved.												
44	Deferred: deferred error. Read-write; Updated-by-hardware. 1=A deferred error was created. A deferred error is the result of an uncorrectable data error which did not immediately cause a processor exception; the data is poisoned and an exception is deferred until the data is loaded by a core or sent to the NB.												
43	Poison: poison error. Read-write; Updated-by-hardware. 1=The error was the result of attempting to consume poisoned data. This indicator does not apply to MSR0000_0411 [NB Machine Check Status (MC4_STATUS)] .												
42:40	Reserved.												
39:36	Way: cache way in error. Read-write; Updated-by-hardware. Indicates the cache way in error. <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>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 2.13.1.5 [Error Code]). See Table 270 for values.												
15:0	ErrorCode: error code. Read-write; Updated-by-hardware. See 2.13.1.5 [Error Code] .												

Table 269: 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

Table 269: LS Error Descriptions

Error Type	Error Sub-type	Description ¹	CTL ²	EAC ³
Internal Error	IntErrType1	A condition was detected which prohibits the core from continuing execution.	IntErrType1e1	E
	IntErrType2		IntErrType1e2	E

1. CID: core ID. All LS errors are reported to the affected core; see 2.13.1.3.
 2. See MSR0000_0400.
 3. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See 2.13.1.3.

Table 270: LS Error Signatures

Error Type	Error SubType	Error-CodeExt	Error Code						UC	ADDRV	PCC	Deferred	Poison			
			Type	PP	T	RRRR	II/TT	LL								
Line Fill Error	-	00001b	MEM	-	-	DRD	D	LG	1	1	0	0	1			
Data Cache Error	Data array	00000b				DRD	D	L1	0/1	1	0	0	0	0		
	SCB	00011b														
	STQ	00010b														
Tag Error	Tag array	10000b				DRD, DWR, Probe	G	L1	0/1	0/1	0/1	0	0	0		
	STQ	10001b				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	0	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	0	0			
Internal Error	IntErrType1	00001b				GEN	1	GEN	GEN	LG	1	0	0	0	0	
	IntErrType2	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 2.13.1 [Machine Check Architecture]. MSR0000_0000 is an alias of MSR0000_0402.

Bits	Description
63:0	ADDR: Address. See Table 271.

Table 271: 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
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	Linear address[47:12]
		4:0	TLB index
	Locked TLB miss	47:12	Linear address[47:12]

MSR0000_0403 LS Machine Check Miscellaneous (MC0_MISC)

See 2.13.1.7 [Error Thresholding].

Bits	Description
63	Valid. IF (PROC>=OR_C0) THEN IF (MSRC001_0015[McStatusWrEn]) THEN Read-write. ELSE Read-only. ENDIF. ELSE Read-write. ENDIF. Reset: 1. 1=A valid CntP field is present in this register.
62	CntP: counter present. IF (PROC>=OR_C0) THEN IF (MSRC001_0015[McStatusWrEn]) THEN Read-write. ELSE Read-only. ENDIF. ELSE Read-write. ENDIF. Reset: 1. 1=A valid threshold counter is present.
61:52	Reserved.
51	CntEn: counter enable. Read-write. Reset: 0. 1=Count thresholding errors; see 2.13.1.7 [Error Thresholding].
50:49	Reserved.
48	Ovrflw: overflow. Read-write; Set-by-hardware. Cold reset: 0. Set by hardware when ErrCnt transitions from FFEh to FFFh. When this field is set, ErrCnt no longer increments.

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). 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 (a write value of FFFh) is not supported.
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 2.13.1 [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.
20	DFIFOE: decoder FIFO parity error.
19	DPDBE: decoder predecode buffer parity error.
18	DEIBP: decoder instruction byte buffer parity error.
17	DEUOPQP: Decoder micro-op queue parity error.
16	DEPRP: microcode patch buffer 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	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	Unused.

MSR0000_0405 IF Machine Check Status (MC1_STATUS)

Cold reset: 0. See 2.13.1 [Machine Check Architecture]. See MSRC001_0015[McStatusWrEn]. Table 272

describes each error type. [Table 273](#) 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	Deferred: deferred error. Read-only. See: MSR0000_0401 [Deferred].								
43	Poison: poison error. See: MSR0000_0401 [Poison].								
42:40	Reserved.								
39:36	Way: cache way in error. Read-write; Updated-by-hardware. Indicates the cache way in error. <table border="1" data-bbox="256 856 649 997"> <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>Fh-2h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	0h	Way 0	1h	Way 1	Fh-2h	Reserved
Bits	Description								
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]. See Table 273 for values.								
15:0	ErrorCode: error code. Read-write; Updated-by-hardware. See 2.13.1.5 [Error Code] .								

Table 272: 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 272: 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. This error may cause subsequent errors related to the entry, but the effect can be contained to the running process.		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	
Instruction cache read error	Microcode Patch Buffer	Parity error in the 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 RAM. If a hard fault is not found, the error was likely a transient upset and the RAM is not broken. This error can also be caused by an error in the microcode patch region of the CC6 save area if ECC is not enabled.	DEPRP	A	E

Table 272: IF Error Descriptions

Error Type	Error Sub-type	Description	CTL ³	CID ²	EAC ¹
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
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

1. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See [2.13.1.3](#).
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 [2.13.1.3](#).
3. See [MSR0000_0404](#).

Table 273: IF Error Signatures

Error Type	Error Sub-type	Error-Code-Ext	Error Code						UC	ADDRV	PCC	Deferred	Poison	
			Type	PP	T	RRRR	II/TT	LL						
Line Fill Error	-	00000b	MEM	-	-	IRD	I	L2	1	1	0	0	1	
Instruction Cache Read Error	IC data load parity	00001b		-	-	IRD	I	L1	0	1	0	0	0	
	IC valid bit	00010b							0	1	0	0	0	
	Main tag	00011b							0	1	0	0	0	
	Prediction queue	00100b							1	0	0	0	0	
	PFB data/addresses	00101b							0/1	0	0	0	0	
	PFB valid bit	01101b							1	0	0	0	0	
	PFB non-cacheable bit	01010b							0/1	0	0	0	0	
	PFB promotion address error	00111b							1	0	1	0	0	
	Branch status register	00110b							1	0	0	0	0	
	Micro-code Patch Buffer	10000b							LG	1	1	1	0	0
	Decoder micro-op queue	10001b							L1	0/1	1	0	0	0
	Decoder instruction buffer	10010b								0/1	1	0	0	0
	Decoder pre-decode buffer	10011b								0/1	0	0	0	0
	Decoder fetch address FIFO	10100b								0/1	1	0	0	0

Table 273: IF Error Signatures

Error Type	Error Sub-type	Error-Code-Ext	Error Code						UC	ADDRV	PCC	Deferred	Poison
			Type	PP	T	RRRR	II/TT	LL					
Tag Probe	Probe tag error	01000b	MEM	-	-	Probe	I	L1	0	1	0	0	0
	Probe tag valid bit	01001b											
L1 TLB	Parity	00000b	TLB	-	-	-	I	L1	0	1	0	0	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	0	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 2.13.1 [Machine Check Architecture].

Bits	Description
63:0	ADDR. See Table 274.

Table 274: 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
Microcode Patch Buffer	-	[63:4] [3:0]	Reserved Line group index.

Table 274: 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 [2.13.1.7 \[Error Thresholding\]](#).

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 [2.13.1 \[Machine Check Architecture\]](#). See [MSRC001_0046 \[CU Machine Check Control Mask \(MC2_CTL_MASK\)\]](#).

Bits	Description
63:15	Unused.
14	L2TlbPoison: TLB fill poison error from L2.
13	RdData: read data error from NB.
12	L2Tag: L2 cache tag error.
11	L2TlbData: L2 TLB parity error.
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 tag 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 [2.13.1 \[Machine Check Architecture\]](#). See [MSRC001_0015\[McStatusWrEn\]](#). [Table 276](#) describes each error type. [Table 277](#) describes the error codes and status register settings for each error type.

Bits	Description										
63	Val: valid. Read-write; Set-by-hardware. See MSR0000_0401[Val] .										
62	Overflow: error overflow. Read-write; Set-by-hardware. See MSR0000_0401[Overflow] .										
61	UC: error uncorrected. Read-write; Updated-by-hardware. See MSR0000_0401[UC] .										
60	En: error enable. Read-write; Updated-by-hardware. See MSR0000_0401[En] .										
59	MiscV: miscellaneous error register valid. Read-write; Updated-by-hardware. See MSR0000_0401[MiscV] . 1=Valid thresholding in MSR0000_040B .										
58	AddrV: error address valid. Read-write; Updated-by-hardware. See MSR0000_0401[AddrV] .										
57	PCC: processor context corrupt. Read-write; Set-by-hardware. See MSR0000_0401[PCC] .										
56:55	Reserved.										
54:47	Syndrome[7:0]. Read-write; Updated-by-hardware. Syndrome[11:0] = {Syndrome[11:8], Syndrome[7:0]}. The syndrome bits when an ECC error is detected. See Table 277 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	Deferred: deferred error. Read-write; Updated-by-hardware. See MSR0000_0401[Deferred] .										
43	Poison: poison error. Read-write; Updated-by-hardware. See MSR0000_0401[Poison] .										
42:40	Reserved.										
39:36	Way: cache way in error. Read-write; Updated-by-hardware. Indicates the cache way in error. See Table 277 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]. Read-write; Updated-by-hardware. See Syndrome[7:0] .										
23:21	Reserved.										
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. See MSR0000_0401[ErrorCodeExt] . See Table 277 for values.										
15:0	ErrorCode: error code. Read-write; Updated-by-hardware. See 2.13.1.5 [Error Code] .										

Table 275: MBE, SBU, and SBC Definitions

Term	Definition
MBE	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 276: CU Error Descriptions

Error Type	Error Sub-Type	Description	CTL ³	CID ²	EAC ¹
System Read Data	L2Tlb	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 276: 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> MBE, 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> MBE, SBU Sync flood. SBC Invalidate Wcc tag entry (cleans error).	WccAddr	0
WccData		WCC data ECC error. <u>SubCase</u> <u>Error Action⁴</u> MBE, 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 276: 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>MBE,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 supercedes the L2Tag error signature if they both occur for the same L2 tag read.</p> <p><u>SubCase</u> <u>Description</u></p> <p>MBE, 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 2.13.1.3.
2. CID: core ID. A=Error reported to the affected core. 0=Error reported to core 0 of the compute unit; see 2.13.1.3.
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 277: 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				-				
System Read Data	L2Tlb	-	00000b	BUS	SRC	0	RD	MEM/IO	L2	1	1	0	-	-	0	0	0	0
	Prefetch	-	00001b					MEM	L2	0								
	Wcc	-	00010b						L1	1								

Table 277: CU Error Signatures

Error Type	Error Sub-Type	Sub Case	Error-Code-Ext	Error Code						UC	ADDRV	PCC	Syn-drome	Way	CECC	UECC	Deferred	Poison													
				Type	PP	T	RRRR	H/TT	LL																						
L2 Cache	FillEcc	MBE	00100b	MEM	-	-	DRD	D	See ¹	0	1	0	[8:0]	See ¹	0	1	1	0													
		SBU													1	0	0	0													
		SBC															0	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	MBE	01001b	MEM	-	-	DWR	D	L1	1	1	1	[11:0]	[1:0]	0	1	0	0													
		SBU													1	0															
		SBC													0	0															
	WccData	MBE	01010b	0	0	[8:0]	0	1	1	0																					
		SBU					1	0	1	0																					
		SBC							0	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													
		MBE													0	1															
		SBU													1	0															
		SBC															0														
Tag	L2Tag	MBE	10000b	MEM	-	-	GEN	G	L2	1	1	1	[7:0]	[3:0]	0	1	0	0													
		SBU													1	0															
		SBC													0	0															
		Hard													10001b	1	1	-	0	0											
	L2Tag MH	-	10010b			1	1	1	-	-	0	0	0	0																	
	XabAddr	-	10011b						-	-																					
	PrbAddr	-	10100b																												

1. LL and Way are specified in [Table 276](#).

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 [2.13.1 \[Machine Check Architecture\]](#).

Bits	Description
63:0	ADDR. See Table 278 .

Table 278: CU Address Register

Error Type	Error Sub-Type	SubCase	Description	
System Read Data Error	TLB	-	<u>Bit</u> [63:48]	<u>Description</u> Reserved.
			[47:6]	PhysAddr[47:6]
TLB	TlbPar	-	<u>Bit</u> [63:7]	<u>Description</u> Reserved.
			[6:0]	Index[6:0]
L2 Cache	FillErr	-	<u>Bit</u> [63:48]	<u>Description</u> Reserved.
			[47:6]	PhysAddr[47:6]
L2 Cache	Prefetch	-	<u>Bit</u> [63:5]	<u>Description</u> Reserved.
			[4:0]	Prefetch FIFO read pointer.
L2 Cache	FillEcc	All	<u>Bit</u> [63:48]	<u>Description</u> Reserved.
			[47:3]	PhysAddr[47:3]. PhysAddr[16:6] is the L2 index and dependent on cache size ¹ . <u>Bit</u> <u>Description</u> [16:6] 2 MB [15:6] 1 MB ([16] Reserved)
L2 Cache	FillPar	L2->LS,TLB	<u>Bit</u> [63:48]	<u>Description</u> Reserved.
			[47:3]	PhysAddr[47:3].
L2 Cache	WccTag	All	[2:0]	Reserved.
			<u>Bit</u> [63:10]	<u>Description</u> Reserved.
L2 Cache	WccData	All	[9:6]	Index[9:6]
			[5:0]	Reserved.
L2 Cache	WcbData	-	<u>Bit</u> [63:48]	<u>Description</u> Reserved.
			[47:3]	PhysAddr[47:3]
L2 Cache	WcbData	-	[2]	Reserved.
			[1:0]	Index[1:0]

Table 278: 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)
			[5:0]	[16:6]: 2 MB [15:6]: 1 MB ([16] Reserved) Reserved.
XabAddr	-	<u>Bits</u>	<u>Description</u>	
PrbAddr	-	[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 [2.13.1.7 \[Error Thresholding\]](#).

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; **Per-node**; **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]. See D18F3x44 [MCA NB Configuration]. See 2.13.1 [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.										
29:28	Unused.										
27	TblWlkDatErrEn: table walk data error enable. 1=Enables MCA reporting of uncorrectable errors in returned data from a GART table walk.										
26	NbArrayParEn: northbridge array parity error reporting enable. 1=Enables reporting of parity errors in the NB arrays.										
25	McaUsPwDatErrEn: MCA upstream data error enable. 1=Enables MCA reporting of upstream posted writes in which the link error bits indicate a data error.										
24	SyncPktEn[3]: link sync packet error reporting enable for link 3. See: SyncPktEn[2:0].										
23	CrcErrEn[3]: link CRC error reporting enable for link 3. See: CrcErrEn[2:0].										
22:19	RtryHtEn: link retry reporting enable. 1=Enables MCA reporting of retries on the link.										
	<table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[3]</td> <td>Link 3</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[0]</td> <td>Link 0</td> </tr> </tbody> </table>	Bit	Description	[3]	Link 3	[2]	Link 2	[1]	Link 1	[0]	Link 0
Bit	Description										
[3]	Link 3										
[2]	Link 2										
[1]	Link 1										
[0]	Link 0										

18	DramParEn: DRAM parity error reporting enable. 1=Enables MCA reporting of parity errors on the DRAM address or control signals.										
17	HtDataEn: link data error reporting enable. 1=Enables MCA reporting of packets with data errors detected on links.										
16	ProtEn: protocol error reporting enable. 1=Enables MCA reporting of protocol errors detected on links or in the L3 cache. When possible, this enable should be cleared before initiating a warm reset to avoid logging spurious errors due to RESET_L signal skew.										
15	L3ArrayUCEn: L3 cache array uncorrectable error reporting enable. 1=Enables MCA reporting of uncorrectable errors in the L3 cache arrays.										
14	L3ArrayCorEn: L3 cache array correctable error reporting enable. 1=Enables MCA reporting of correctable errors in the L3 cache arrays.										
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	GartTblWkEn: GART table walk error reporting enable. 1=Enables MCA reporting of GART cache table walks which encounter a GART PTE entry which is invalid.										
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.										
7:5	<p>SyncPktEn[2:0]: link sync packet error reporting enable for links [2:0]. SyncPktEn[3:0] = {SyncPktEn[3], SyncPktEn[2:0]}. 1=Enables MCA reporting of 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 this bit.</p> <table border="1"> <thead> <tr> <th><u>Bit</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>[3]</td> <td>Link 3</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[0]</td> <td>Link 0</td> </tr> </tbody> </table>	<u>Bit</u>	<u>Description</u>	[3]	Link 3	[2]	Link 2	[1]	Link 1	[0]	Link 0
<u>Bit</u>	<u>Description</u>										
[3]	Link 3										
[2]	Link 2										
[1]	Link 1										
[0]	Link 0										

4:2	CrcErrEn[2:0]: link CRC error reporting enable for links [2:0]. CrcErrEn[3:0] = {CrcErrEn[3], CrcErrEn[2:0]}. 1=Enables MCA reporting of CRC errors detected on link (see the description of CRC Error in Table 279). The NB floods its outgoing links with sync packets after detecting a CRC error on an incoming link independent of the state of this bit. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[3]</td> <td>Link 3</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[0]</td> <td>Link 0</td> </tr> </tbody> </table>	Bit	Description	[3]	Link 3	[2]	Link 2	[1]	Link 1	[0]	Link 0
Bit	Description										
[3]	Link 3										
[2]	Link 2										
[1]	Link 1										
[0]	Link 0										
1	UECCEn: uncorrectable ECC error reporting enable. 1=Enables MCA reporting of DRAM uncorrectable ECC errors which are detected in the NB. In some cases data may be forwarded to the core prior to checking ECC in which case the check takes place in one of the other error reporting banks.										
0	CECCEn: correctable ECC error reporting enable. 1=Enables MCA reporting of DRAM correctable ECC errors which are detected in the NB.										

MSR0000_0411 NB Machine Check Status (MC4_STATUS)

Per-node; Not-same-for-all. Cold reset: x. See [2.13.1 \[Machine Check Architecture\]](#). See [MSRC001_0015\[McStatusWrEn\]](#). [Table 279](#) describes each error type. [Table 280](#) and [Table 281](#) 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. Read-write; Set-by-hardware. See MSR0000_0401[Val] .
62	Overflow: error overflow. Read-write; Set-by-hardware. See MSR0000_0401[Overflow] .
61	UC: error uncorrected. Read-write; Updated-by-hardware. See MSR0000_0401[UC] .
60	En: error enable. Read-write; Updated-by-hardware. See MSR0000_0401[En] .
59	MiscV: miscellaneous error register valid. Read-write; Updated-by-hardware. See MSR0000_0401[MiscV] . 1=Valid thresholding in MSR0000_0413 , MSRC000_0408 , or MSRC000_0409 .
58	AddrV: error address valid. Read-write; Updated-by-hardware. See MSR0000_0401[AddrV] .
57	PCC: processor context corrupt. Read-write; Set-by-hardware. See MSR0000_0401[PCC] .
56	ErrCoreIdVal: error core ID is valid. Read-write; Set-by-hardware. 1=The ErrCoreId field is valid.
55	Reserved.
54:47	Syndrome[7:0]. Read-write; Updated-by-hardware. Syndrome[15:0] = {Syndrome[15:8], Syndrome[7:0]}. The syndrome bits when an ECC error is detected. See Table 281 Valid Syndrome column for which bits are valid for each error.
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	Reserved.

43:42	McaStatSubCache: L3 subcache in error. Read-write; Updated-by-hardware. Indicates the number of the L3 subcache associated with the error. This field is only valid when an L3 error is recorded.																		
41	SubLink: sublink or DRAM channel. 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. For DRAM parity errors, this bit indicates which channel the error was associated with. 0=DCT A; Sublink [7:0]. 1=DCT B; Sublink [15:8].																		
40	Scrub: error detected on a scrub. Read-write; Set-by-hardware.																		
39:36	<p>Link. Read-write; Set-by-hardware.</p> <ul style="list-style-type: none"> For errors associated with a link, this field indicates which link was associated with the error. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[3]</td> <td>Link 3</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[0]</td> <td>Link 0</td> </tr> </tbody> </table> <ul style="list-style-type: none"> For L3 cache errors, this field indicates the L3 way in error, and McaStatSubCache contains sub-cache number. For DRAM parity errors, this field indicates which DCT detected an error. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[3:2]</td> <td>Reserved</td> </tr> <tr> <td>[1]</td> <td>DCT 1</td> </tr> <tr> <td>[0]</td> <td>DCT 0</td> </tr> </tbody> </table>	Bit	Description	[3]	Link 3	[2]	Link 2	[1]	Link 1	[0]	Link 0	Bit	Description	[3:2]	Reserved	[1]	DCT 1	[0]	DCT 0
Bit	Description																		
[3]	Link 3																		
[2]	Link 2																		
[1]	Link 1																		
[0]	Link 0																		
Bit	Description																		
[3:2]	Reserved																		
[1]	DCT 1																		
[0]	DCT 0																		
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	Syndrome[15:8]. Read-write; Updated-by-hardware. See Syndrome[7:0].																		
23:21	Reserved.																		
20:16	ErrorCodeExt: extended error code. Read-write; Updated-by-hardware. See MSR0000_0401[ErrorCodeExt]. See Table 280 for values.																		
15:0	ErrorCode: error code. Read-write; Updated-by-hardware. See 2.13.1.5 [Error Code].																		

Table 279: NB Error Descriptions

Error Type	Description	CTL ¹	ETG ²	EAC
CRC Error	<p>CRC error detected on link. If the link is in retry mode, this may indicate excessive link reconnect failures; see D18F0x[E4,C4,A4,84][CrcErr, LinkFail, CrcFloodEn]. This error can also be seen as the result of link training failures.</p> <p>The NB floods its outgoing links with sync packets after detecting a CRC error on an incoming link independent of the state of the control bits.</p>	CrcErrEn	L	D
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

Table 279: NB Error Descriptions

Error Type	Description	CTL ¹	ETG ²	EAC
Master Abort	Master abort seen as result of link operation. Reasons for this error include requests to non-existent addresses, and requesting extended addresses while extended mode disabled (see D18F0x[E4,C4,A4,84][Addr64BitEn]). 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
GART Error	GART cache table walk encountered a GART PTE entry which was invalid.	GartTblWkEn	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
ECC Error	DRAM ECC error detected.	CECCEn, UECCEn	D	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.)	HtDataEn, McaUsPwDatErrEn, D18F3x180[DeferDatErrNcHtMcaEn]	L	D
Protocol Error	Protocol error detected by link, L3, or probe filter. These errors are distinguished from each other by the value in MSR0000_0412[ErrAddr] . See Table 283 . 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. For link protocol errors, ensure that the error is not due to failure or reset at the far end of link or from transmission corruption, indicated by CRC error. The enable for this error should be cleared before initiating a warm reset to avoid logging spurious errors due to RESET_L signal skew.	ProtEn	L ³	D
NB Array Error	A parity error was detected in the NB internal arrays.	NbArrayParEn	-	D
DRAM Parity Error	A parity error was detected on the DRAM address or control signals.	DramParEn	D	D

Table 279: NB Error Descriptions

Error Type	Description	CTL ¹	ETG ²	EAC
Link Retry	A transmission error occurred on the link; the IO link Error Retry Protocol is executed. Retry may have been initiated by either end of the link.	RtryHtEn	L	D
GART Table Walk Data Error	An uncorrectable error was found in data returned from a GART table walk.	TblWlkDatErrEn	L	D
L3 Cache Data Error	ECC error detected in L3 cache data.	L3ArrayCorEn, L3ArrayUCEn	C	D
L3 Cache Tag Error	Error detected in L3 cache tag. The subcache, index, and way are logged. See Table 281 footnotes.		C	D
L3 Cache LRU Error	Error detected in LRU parity bits. This is a non-fatal error which has no impact on any program execution; LRU state is reset. The cache index is captured for thresholding purposes.	-	C	D
Probe Filter Error	An ECC error was detected in the probe filter directory. See 2.9.4.1.1 [Probe Filter Errors] .	CECCEn	C	D
Compute Unit Data Error	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: <ul style="list-style-type: none"> • DRAM • APIC • Configuration space (IO and MMIO) For these errors, sync flood will occur if D18F3x180[CpuLeakErrEn] is set.	McaCpuDatErrEn	-	D

1. See [MSR0000_0410](#).

2. ETG: error threshold group. L=Link. D=DRAM. C=L3 Cache. See [2.13.1.7 \[Error Thresholding\]](#).

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 [2.13.1.3](#).

The NB is capable of reporting the following errors:

Table 280: NB Error Signatures, Part 1

Error Type	ErrorCode-Ext	Error Code					
		Type	PP	T	RRRR	II/TT	LL
Reserved	00000b	-	-	-	-	-	-
CRC Error	00001b	BUS	OBS	0	GEN	GEN	LG
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
GART Error	00101b	TLB	-	-	-	GEN	LG
RMW Error	00110b	BUS	OBS	0	GEN	IO	LG
WDT Error	00111b	BUS	GEN	1	GEN	GEN	LG
ECC Error	01000b	BUS	SRC/RES	0	RD/WR	MEM	LG

Table 280: NB Error Signatures, Part 1

Error Type	ErrorCode-Ext	Error Code					
		Type	PP	T	RRRR	II/TT	LL
Link Data Error	01010b	BUS	SRC/OBS	0	RD/WR/DWR	MEM/IO	LG
Protocol Error	01011b	BUS	OBS	0	GEN	GEN	LG
NB Array Error	01100b	BUS	OBS	0	GEN	GEN	LG
DRAM Parity Error	01101b	BUS	OBS	0	GEN	MEM	LG
Link Retry	01110b	BUS	OBS	0	GEN	GEN	LG
GART Table Walk Data Error	01111b	TLB	-	-	-	GEN	LG
L3 Cache Data Error	11100b	MEM	-	-	RD/Evict /Probe /GEN	GEN	LG
L3 Cache Tag Error	11101b						
L3 Cache LRU Error	11110b						
Probe Filter Error	11111b						
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 281: NB Error Signatures, Part 2

Error Type	UC	AddrV	PCC	Syndrome Valid	CECC	UECC	Scrub	Link	Err CoreId
CRC Error	1	0	1	-	0	0	0	Y	-
Sync Error	1	0	1	-	0	0	0	Y	-
Mst Abort	1	1	Core ¹²	-	0	0	0	Y	Y
Tgt Abort	1	1	Core ¹²	-	0	0	0	Y	Y
GART Error	1	1	Core ¹²	-	0	0	0	-	Y
RMW Error	1	1	0	-	0	0	0	Y	-
WDT Error	1	0 ¹	1	-	0	0	0	-	-
ECC Error	MS ⁸	1	MS ⁸ & Core ¹²	15:0	~MS ⁸	MS ⁸	1/0	-	-
Link Data Error	1	1	0	-	0	0	0	Y	-
Protocol Error	1	1/0 ²	1	-	0	0	0	Y	-
NB Array Error	1	1 ⁴	1	-	0	0	0	-	-
DRAM Parity Error	1	0	1	-	0	0	0	-	-
Link Retry ⁹	0	0	0	-	0	0	0	Y	-
GART Table Walk Data Error	1	1	0	-	0	0	0	-	-

Table 281: NB Error Signatures, Part 2

Error Type	UC	AddrV	PCC	Syndrome Valid	CECC	UECC	Scrub	Link	Err CoreId
L3 Cache Data Error	MB ¹⁰	1 ⁵	If UC	15:0	SB ¹¹	MB ¹⁰	0/1	Y ⁶	Y ⁷
L3 Cache Tag Error									
L3 Cache LRU Error	0	1 ⁵	0	-	0	0	0	-	-
Probe Filter Error	0	1 ⁵	0	15:0	SB ¹¹	MB ¹⁰	0/1	Y ⁶	Y ⁷
Compute Unit Data Error	1	0	1	-	0	0	0	-	Y

1. See [Table 287: \[NB Address Register for Watchdog Timer Errors\]](#)
2. See [Table 283 \[NB Address Register for Protocol Errors\]](#)
3. Link identified only if link protocol error. See entry in [Table 279](#).
4. See [Table 285: \[NB Address Register for NB Array Errors\]](#)
5. See [Table 286: \[NB Address Register for L3 Array Errors\]](#)
6. This field contains the L3 way in error. [D18F3x4C\[McaStatSubCache\]](#) contains the subcache number.
7. Depends on Memory Transaction Type ([Table 123](#)); valid if non-zero.
8. MS: multi-symbol. 1=Multi-symbol. 0=Not multi-symbol. See [2.13.2 \[DRAM Considerations for ECC\]](#) for information on symbol size.
9. Retries initiated by either side of the link are logged.
10. MB: multi-bit error. 1=Multi-bit error.
11. SB: single bit error. 1=Single bit error.
12. Core: source is core. 1=Source is core. 0=Source is not core.

MSR0000_0412 NB Machine Check Address (MC4_ADDR)

Read-write; Per-node; Not-same-for-all. Cold reset: x. See [2.13.1 \[Machine Check Architecture\]](#).

[MSR0000_0412\[31:0\]](#) is an alias of [D18F3x50](#). [MSR0000_0412\[63:32\]](#) is an alias of [D18F3x54](#).

[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 283](#).

- NB array errors indicate the array in error, and are formatted according to [Table 285](#).
- L3 array errors store the physical address which caused the error, and are formatted according to [Table 286](#).
- 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 282](#). If [D18F3x4C\[AddrV\]](#) is not indicated, errors are formatted according to [Table 287](#).
- All other NB errors which indicate [D18F3x4C\[AddrV\]](#) are formatted according to [Table 282](#).

Table 282: NB Address Register Default Encoding

ErrAddr	Description
47:1	Physical address bits 47:1.

Table 283: 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 284 .

Table 284: 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	L3: Request gets multiple hits in L3
10001b	L3: Probe access gets multiple hits in L3
10010b	L3: Request queue overflow
10011b	L3: WrVicBlk hit incompatible L3 state

Table 284: Protocol Error Type

Link Protocol Error Type	Description
10100b	L3: CIVicBlk hit incompatible L3 state
10101b-10111b	Reserved
11000b	PF: Directed probe miss
11001b	PF: Directed probe clean hit
11010b	PF: VicBlkM hit inconsistent directory state 'O S'
11011b	PF: VicBlkE hit inconsistent directory state 'O S'
11100b	PF: Reserved
11101b	PF: L3 lookup response without a matching PFQ entry
11110b	PF: L3 update data read request without a matching PFQ entry
11111b	PF: Reserved

Table 285: 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	LPB1: Link 1 packet buffer
	00111b	LPB2: Link 2 packet buffer
	01000b	LPB3: Link 3 packet buffer
	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	LPS0: Link 0 packet state buffer
	10001b	LPS1: Link 1 packet state buffer
	10010b	LPS2: Link 2 packet state buffer
	10011b	LPS3: Link 3 packet state buffer
	10100b	RHB0: Link 0 retry history buffer
	10101b	RHB1: Link 1 retry history buffer
	10110b	RHB2: Link 2 retry history buffer
	10111b	RHB3: Link 3 retry history buffer
	11000b	SRIMCTRTE: SRI/MCT extended routing table
	11001b	LN0LN1RTE: Link 0/1 extended routing table
	11010b	LN2LN3RTE: Link 2/3 extended routing table
	11011b	GART: GART array
	11100b	Reserved
	11101b	TCB: TCB array.
	11110b	Reserved
	11111b	Reserved

Table 286: NB Address Register for L3 Array Errors

Error Type	Memory Transaction Type (RRRR; Table 123)	ErrAddr bits ¹	Description
Data Error	RD, Evict	47:2	Physical address
	Probe	47:6	Cache line address requested
		5:4	Index of the critical OW within the cache line
		3:2	Index of the OW in error within the cache line (normally occupies bits 5:4 in physical address)
	GEN	16:6	Cache index
Tag or LRU error	RD, Evict, Probe	47:6	Physical address contained in the tag (may not match the address requested)
	GEN	16:6	Cache index

1. The physical address includes the cache index in bits 16:6.

Table 287: 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"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>07h-00h</td> <td>CoreId</td> </tr> <tr> <td>FFh-08h</td> <td>Reserved</td> </tr> </tbody> </table>	Bits	Description	07h-00h	CoreId	FFh-08h	Reserved												
Bits	Description																		
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.																		
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>GART Table Walker</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	GART Table Walker	10b	Memory Controller	11b	Host								
Bits	Description																		
00b	Core																		
01b	GART Table Walker																		
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																		

Table 287: NB Address Register for Watchdog Timer Errors

Bits	Description
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)

Per-node; Not-same-for-all. [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](#) and [MSRC000_0409](#). See [D18F3x44\[NbMcaToMstCpuEn\]](#). [MSR0000_0413\[63:32\]](#) is an alias of [D18F3x160](#).

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. Read-write. Reset: 0h. BIOS: 1h. 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. Read-write. Reset: 0. 1=Count thresholding errors; see 2.13.1.7 [Error Thresholding] .										
50:49	IntType: interrupt type. Read-write. 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. Read-write; Set-by-hardware. Cold reset: 0. Set by hardware when ErrCnt transitions from FFEh to FFFh. When this field is set, ErrCnt no longer increments. When this bit is set, the interrupt selected by the IntType field is generated.										
47:44	Reserved.										
43:32	ErrCnt: error counter. Read-write; Updated-by-hardware. 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 (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 and MSRC000_0409 .										
23:0	Reserved.										

MSR0000_0414 EX Machine Check Control (MC5_CTL)

Read-write. Reset: 0000_0000_0000_0000h. See [2.13.1 \[Machine Check Architecture\]](#). 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 error.
10	AG1PRF: physical register file AG1 port parity error.
9	AG0PRF: physical register file AG0 port parity error.
8	EX1PRF: physical register file EX1 port parity error.
7	EX0PRF: physical register file EX0 port parity error.
6	MAP: mapper checkpoint array parity error.
5	RETDISP: retire dispatch queue parity error.
4	IDF: IDRF array parity error.
3	PLDEX: EX payload array parity error.
2	PLDAG: AG payload array parity error.
1	PICWAK: wakeup array dest tag parity error.
0	Unused.

MSR0000_0415 EX Machine Check Status (MC5_STATUS)

Cold reset: 0. See 2.13.1 [Machine Check Architecture]. See MSRC001_0015[McStatusWrEn]. Table 288 describes each error type. Table 289 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]. See Table 289 for values.
15:0	ErrorCode: error code. Read-write; Updated-by-hardware. See 2.13.1.5 [Error Code].

Table 288: 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; 2.13.1.3.1 [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; 2.13.1.3.1 [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"> 1. See MSR0000_0414. 2. CID: core ID. All EX errors are reported to the affected core; see 2.13.1.3. 3. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See 2.13.1.3. 				

Table 289: EX Error Signatures

Error Type	Error Sub-Type	Error Code Ext	Error Code						UC	ADDRV	PCC
			Type	PP	T	RRRR	II	LL			
WDT error	-	00h	BUS	GEN	1	GEN	GEN	LG	1	1	1
Internal error	Wakeup array dest tag parity	01h			0				1/0	0	0
	AG payload array parity	02h							1/0	1	0
	EX payload array parity	03h							1/0	1	0
	IDRF array parity	04h							1/0	1	0
	Retire dispatch queue parity	05h							1 ¹	1	0
	Mapper checkpoint array parity	06h							1/0 ²	1	0
	EX0PRF parity	07h							1/0	0	0
	EX1PRF parity	08h							1/0	0	0
	AG0PRF parity	09h							1/0	0	0
	AG1PRF parity	0Ah							1/0	0	0
	Flag register file parity	0Bh							1/0	0	0
	DE error	0Ch							1/0	0	0

1. Causes shutdown.
2. Causes 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 2.13.1 [Machine Check Architecture]. The register format depends on the type of error being logged:

Bits	Description
63:0	ADDR . See Table 290.

Table 290: EX Address Register

Error Type	Error Sub-Type	Bits	Description
WDT	- (ErrorCodeExt=00h)	63:48	Reserved
		47:0	LogAddr[47:0] . Logical address of the next instruction after the last instruction retired.
Internal	AG payload array parity (ErrorCodeExt=02h), EX payload array parity (ErrorCodeExt=03h), IDRF array parity (ErrorCodeExt=04h)	63:6	Reserved
		5:0	SchedulerQID .

Table 290: EX Address Register

Error Type	Error Sub-Type	Bits	Description
Internal	Retire dispatch queue parity (ErrorCode-Ext=05h)	63:7	Reserved
		6:0	RetirementID.
Internal	Mapper checkpoint array parity (ErrorCode-Ext=06h)	63:6	Reserved
		5:0	CheckpointID.

MSR0000_0417 EX Machine Check Miscellaneous (MC5_MISC)

See 2.13.1.7 [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 2.13.1 [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)

See 2.13.1 [Machine Check Architecture]. See [MSRC001_0015\[McStatusWrEn\]](#). [Table 291](#) describes each error type. [Table 292](#) 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. Value: 0. See MSR0000_0401[MiscV] .
58	AddrV: error address valid. Read-only. Value: 0. 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. Cold reset: 0. See MSR0000_0401[ErrorCodeExt] . See Table 292 for values.
15:0	ErrorCode: error code. Read-write; Updated-by-hardware. Cold reset: 0. See 2.13.1.5 [Error Code] .

Table 291: 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

1. EAC: D=Error action taken if detected. E=Error action taken if MCA bank enabled. See [2.13.1.3](#).

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 [2.13.1.3](#).

4. See [MSR0000_0418](#).

Table 292: 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 2.13.1 [Machine Check Architecture].

Bits	Description
63:0	Reserved.

MSR0000_041B FP Machine Check Miscellaneous (MC6_MISC)

Reset: 0000_0000_0000_0000h. Read-only. See 2.13.1 [Machine Check Architecture].

Bits	Description
63:0	Reserved.

3.12 MSRs - MSRC000_0xxx

MSRC000_0080 Extended Feature Enable (EFER)

SKINIT Execution: 0000_0000_0000_0000h.

Bits	Description
63:16	MBZ.
15	IF (PROC >= OR_C0) THEN TCE: translation cache extension enable. Read-write. 1=Translation cache extension is enabled. See CPUID Fn8000_0001_EDX[TCE] . ELSE MBZ. ENDIF.
15	MBZ.
14	FFXSE: fast FXSAVE/FRSTOR enable. Read-write. Reset: 0. 1=Enables the fast FXSAVE/FRSTOR mechanism. Support for this field indicated by CPUID Fn8000_0001_EDX[FFXSR] . 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. Reset: 0. 1=Enables the long mode segment limit check mechanism.
12	SVME: secure virtual machine (SVM) enable. IF (MSRC001_0114 [SvmeDisable]==1) THEN MBZ. ELSE Read-write. ENDIF. Reset: 0. 1=SVM features are enabled.
11	NXE: no-execute page enable. Read-write. Reset: 0. 1=The no-execute page protection feature is enabled. Support for this field indicated by CPUID Fn8000_0001_EDX[NX] .
10	LMA: long mode active. Read-only; Not-same-for-all. Reset: 0. 1=Indicates that long mode is active. Support for this field indicated by CPUID Fn8000_0001_EDX[LM] .
9	MBZ.
8	LME: long mode enable. Read-write; Not-same-for-all. Reset: 0. 1=Long mode is enabled. Support for this field indicated by CPUID Fn8000_0001_EDX[LM] .
7:1	RAZ.
0	SYSCALL: system call extension enable. Read-write. Reset: 0. 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)

Bits	Description
63:32	RAZ.
31:0	Mask: SYSCALL flag mask. Read-write. Reset: 0000_0000h. This register holds the EFLAGS mask used by the SYSCALL instruction. 1=Clear the corresponding EFLAGS bit when executing the SYSCALL instruction.

IF (PROC>=OR_C0) THEN

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. A read of this MSR in guest mode is affected by MSRC000_0104 [Time Stamp Counter Ratio (TscRateMsr)]. This field uses software P-state numbering. See MSRC001_0015[EffFreqCntMwait], 2.5.6 [Effective Frequency Interface], and 2.5.2.1.2.1 [Software P-state Numbering]. This register is not affected by writes to MSR0000_00E7.

ENDIF.

IF (PROC>=OR_C0) THEN

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.

ENDIF.

MSRC000_0100 FS Base (FS_BASE)

Updated-by-hardware. 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)

Updated-by-hardware. Reset: 0000_0000_0000_0000h.

Bits	Description
63:0	GSBase: expanded GS segment base. Read-write. 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)

Updated-by-hardware. 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)

Updated-by-hardware. 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)

[MSRC000_0104 \[Time Stamp Counter Ratio \(TscRateMsr\)\]](#) allows the hypervisor to control the guest's view of the Time Stamp Counter. It provides a multiplier that scales the value returned when [MSR0000_0010\[TSC\]](#) and [MSR0000_00E7\[MPERF\]](#) are read by a guest running under virtualization. This allows the hypervisor to provide a consistent TSC and MPERF rate for a guest process when moving that process between cores that have a differing P0 rate. The TSC Ratio MSR does not affect the value read from the TSC and MPERF MSRs when read when in host mode or when virtualization is not being used or when accessed by code executed in system management mode (SMM) unless the SMM code is executed within a guest container. The TSC Ratio value does not affect the rate of the underlying TSC and MPERF counters, or the value that gets written to the TSC and MPERF MSRs counters on a write by either the host or the guest. The TSC Ratio MSR contains a fixed-point number in 8.32 format, which is 8 bits of integer and 32 bits of fraction. This number is the ratio of the desired P0 frequency to the P0 frequency of the core. The reset value of the TSC Ratio MSR is 1.0, which results in a guest frequency matches the core P0 frequency.

Bits	Description
63:40	MBZ.
39:32	TscRateMsrInt: time stamp counter rate integer. Read-write. Reset: 01h. Specifies the integer part of the MSR TSC ratio value.
31:0	TscRateMsrFrac: time stamp counter rate fraction. Read-write. Reset: 0000_0000h. Specifies the fractional part of the MSR TSC ratio value.

MSRC000_0105 Lightweight Profile Configuration (LWP_CFG)

Bits	Description
63:48	Reserved.
47:40	LwpVector: threshold interrupt vector. Read-write. Reset: 0. Interrupt vector number used by LWP Threshold interrupts. Must be provided if LwpInt is set to 1.
39:32	LwpCoreId: core ID. Read-write. Reset: 0. 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. Reset: 0. 1=Enable LWP to interrupt on threshold overflow. See CPUID Fn8000_001C_EAX[LwpInt] .
30:7	MBZ.
6	LwpRNH: core reference clocks not halted event support. MBZ. Reset: 0. 1=Enable LWP to count core reference clocks not halted. See CPUID Fn8000_001C_EAX[LwpRNH] .
5	LwpCNH: core clocks not halted event support. MBZ. Reset: 0. 1=Enable LWP to count core clocks not halted. See CPUID Fn8000_001C_EAX[LwpCNH] .
4	LwpDME: DC miss event support. MBZ. Reset: 0. 1=Enable LWP to count DC misses. See CPUID Fn8000_001C_EAX[LwpDME] .
3	LwpBRE: branch retired event support. Read-write. Reset: 0. 1=Enable LWP to count branches retired. See CPUID Fn8000_001C_EAX[LwpBRE] .
2	LwpIRE: instructions retired event support. Read-write. Reset: 0. 1=Enable LWP to count instructions retired. See CPUID Fn8000_001C_EAX[LwpIRE] .
1	LwpVAL: LWPVAL instruction support. Read-write. Reset: 0. 1=LWPVAL instruction is enabled. See CPUID Fn8000_001C_EAX[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 will cause 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)

Per-node; Not-same-for-all. MSRC000_0408 is associated with the link error type. MSRC000_0408 and MSRC000_0409 are the block of extended NB machine check miscellaneous registers. MSRC000_0408[63:32] is an alias of D18F3x168.

Bits	Description
63	Valid. See: MSR0000_0413 [Valid].
62	CntP: counter present. See: MSR0000_0413 [CntP].
61	Locked. See: MSR0000_0413 [Locked].
60:56	Reserved.
55:52	LvtOffset: LVT offset. See: MSR0000_0413 [LvtOffset]. BIOS: 1h.
51	CntEn: counter enable. See: MSR0000_0413 [CntEn].
50:49	IntType: interrupt type. See: MSR0000_0413 [IntType].
48	Ovrflw: overflow. See: MSR0000_0413 [Ovrflw].
47:44	Reserved.
43:32	ErrCnt: error counter. See: MSR0000_0413 [ErrCnt].
31:24	BlkPtr: Block pointer for additional MISC registers. See: MSR0000_0413 [BlkPtr].
23:0	Reserved.

MSRC000_0409 Machine Check Misc 4 (L3 Thresholding) 2 (MC4_MISC2)

Per-node; Not-same-for-all. MSRC000_0409 is associated with the L3 cache error type. See [MSRC000_0408](#). MSRC000_0409[63:32] is an alias of [D18F3x170](#).

Bits	Description
63	Valid. See: MSR0000_0413 [Valid].
62	CntP: counter present. See: MSR0000_0413 [CntP].
61	Locked. See: MSR0000_0413 [Locked].
60:56	Reserved.
55:52	LvtOffset: LVT offset. See: MSR0000_0413 [LvtOffset]. BIOS: 1h.
51	CntEn: counter enable. See: MSR0000_0413 [CntEn].
50:49	IntType: interrupt type. See: MSR0000_0413 [IntType].
48	Ovrflw: overflow. See: MSR0000_0413 [Ovrflw].
47:44	Reserved.
43:32	ErrCnt: error counter. See: MSR0000_0413 [ErrCnt].
31:24	BlkPtr: Block pointer for additional MISC registers. See: MSR0000_0413 [BlkPtr].
23:0	Reserved.

MSRC000_040[F:A] Reserved

Bits	Description
63:0	RAZ.

3.13 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 293: [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	See: 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 294: [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	See: MSRC001_020[7,5,3,1] .

MSRC001_0010 System Configuration (SYS_CFG)

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.
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]. 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. BIOS: This bit should be set to 1 during BIOS initialization of the fixed MTRRs, then cleared to 0 for operation.
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 RDBlkM command.
15:0	Reserved.

MSRC001_0015 Hardware Configuration (HWCR)

Bits	Description
63:28	Reserved.
27	IF (PROC >= OR_C0) THEN EffFreqReadOnlyLock: read-only effective frequency counter lock. Write-1-only. Reset: 0. BIOS: 1. 1= MSRC000_00E7 and MSRC000_00E8 are read-only. ELSE Reserved. ENDIF.
27	Reserved.
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 2.5.6 [Effective Frequency Interface] .
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 2.5.2.1.1 [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 2.5.2.1.2.1 [Software P-state Numbering] .

23	<p>ForceRdWrSzPrb: force probes for RdSized and WrSized. Read-write. Reset: 0. A read returns a 1 if this field is set on any core of the node.</p> <p>1=Forces probes on read-sized and write-sized transactions:</p> <ul style="list-style-type: none"> • From the core that are not to the GART aperture. • From IO, upstream: <ul style="list-style-type: none"> • Except for display refresh transactions (Isoc=1, PassPW=1, RespPassPW=1, Coherent=0) and • Except operations that are “Isoc=1, PassPW=1, RespPassPW=1, Coherent=1”. If this case is to the GART aperture, the GART entry’s “Coherent” bit is used and not forced to Coherent=1. <p>See 2.9.3.1.1 [DRAM and MMIO Memory Space].</p> <p>Also applies to table walks for IO accesses to GART aperture.</p> <p>ForceRdWrSzPrb==1 overrides D18F3x90[DisGartTblWlkPrb].</p>
22:21	Reserved.
20	<p>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)].</p>
19	Reserved.
18	<p>McStatusWrEn: machine check status write enable. Read-write. Reset: 0. McStatusWrEn can be used to debug machine check exception and interrupt handlers. See 2.13.3 [Error Injection and Simulation]. See 2.13.1 [Machine Check Architecture].</p> <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 (D18F3x160, D18F3x168, and D18F3x170). • 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_0413, MSRC000_0408, MSRC000_0409.
17	<p>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.</p>
16:15	Reserved.
14	<p>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.</p>

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. 1=A link special bus cycle is generated when exiting from the Halt state. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E].
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].
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. See 2.4.8.2.9 [Locking SMM].

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. Some key points to consider:

- The IORRs should not cover the range used for the AGP aperture if the GART logic in the NB is enabled.
- The IORRs should be programmed to cover the AGP aperture if the aperture/GART translation is handled by an IO device (e.g., the chipset).

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 NB Configuration 1 (NB_CFG1)

Read-write; Per-node. 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. See 2.12.8 [Response Ordering].
49:47	Reserved.
46	EnableCf8ExtCfg: enable CF8 extended configuration cycles. Reset: 0. 1=Allows the IO configuration space access method, IOCF8 and IOCFC, 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. BIOS: IF (DataMaskMbType != 1) THEN 1 ELSE 0 ENDIF. 1=Disables DRAM data masking function; all write requests that are less than one cacheline, a DRAM read is performed before writing the data. If x4 DIMMs are present, data masking is disabled regardless of the value of this bit. Data masking is supported in ECC mode; the NB performs a minimum write size of 16B.
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.
30:11	Reserved.
10	DisXdsBypass: disable xbar data scheduler bypass. Reset: 0. 1=The crossbar data scheduler bypass is disabled. This bit should be set in systems containing coherent devices that are not covered by this document.
9	DisRefUseFreeBuf: disable display refresh to use free list buffers. Reset: 0. 1=In non-IFCM disable display refresh requests from using free list buffers and in IFCM disable isochronous requests from using free list buffers.
8:0	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 296](#). 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 295: [Register Mapping](#) for [MSRC001_00\[35:30\]](#)

Register	Function
MSRC001_0030	Characters 7-0
MSRC001_0031	Characters 15-8
MSRC001_0032	Characters 23-16
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 296: [BIOS Recommendations](#) 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)

[Per-node](#).

Bits	Description
63:32	Reserved.
31:0	Alias of D18F3x64 . Uses hardware P-state numbering. See 2.5.2.1.2.2 [Hardware P-state Numbering] . See MSRC001_0072[NumBoostStates] .

MSRC001_0044 LS Machine Check Control Mask (MC0_CTL_MASK)

Read-write. Reset: 0000_0000_0000_0000h. BIOS: 0000_0000_0000_0000h. 1=Inhibit detection of error source. See [2.13.1 \[Machine Check Architecture\]](#). See [MSR0000_0400 \[LS Machine Check Control \(MC0_CTL\)\]](#).

Bits	Description
63:10	Reserved.
9	Reserved.
8	IntErrType2: internal error type 2.
7	SRDE: read data errors. System read data errors on cache fill.
6	LineFillPoison: line fill poison error.
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 error.

MSRC001_0045 IF Machine Check Control Mask (MC1_CTL_MASK)

Read-write; **SharedC**. Reset: 0000_0000_0000_0080h. BIOS: 0000_0000_0000_0080h. See [2.13.1 \[Machine Check Architecture\]](#). 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 byte buffer parity error. See IDP.
17	DEUOPQP: Decoder micro-op queue parity error.
16	DEPRP: microcode patch buffer 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. If IDP is set to 1 then DEIBP must be set to 1.
1:0	Reserved.

MSRC001_0046 CU Machine Check Control Mask (MC2_CTL_MASK)

SharedC. Reset: 0000_0000_0000_0000h. BIOS: 0000_0000_0000_0000h. See 2.13.1 [Machine Check Architecture]. See MSR0000_0408 [CU Machine Check Control (MC2_CTL)].

Bits	Description
63:0	See: MSR0000_0408. The format of MC2_CTL_MASK corresponds to MC2_CTL.

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; Per-node; Not-same-for-all. BIOS: 0000_0000_0000_0400h. The format of MC4_CTL_MASK corresponds to MSR0000_0410 [NB Machine Check Control (MC4_CTL)]. For each defined bit position 1=Disable logging. See 2.13.1 [Machine Check Architecture]. See MSR0000_0410 [NB Machine Check Control (MC4_CTL)].

Bits	Description
63:32	Reserved.
31	McaCpuDatErrEn. Reset: 0.
30	Reserved.
29:28	Reserved.
27	TblWlkDatErrEn. Reset: 0.
26	NbArrayParEn. Reset: 0.
25	McaUsPwDatErrEn. Reset: 0.
24	SyncPktEn[3]. Reset: 0. See: SyncPktEn[2:0].
23	CrcErrEn[3]. Reset: 0. See: CrcErrEn[2:0].

22:19	RtryHtEn. Reset: 0. BIOS is recommended to mask logging HT retries if the OS is not capable of distinguishing that HT retries are normal operation. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[3]</td> <td>Link 3</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[0]</td> <td>Link 0</td> </tr> </tbody> </table>	Bit	Description	[3]	Link 3	[2]	Link 2	[1]	Link 1	[0]	Link 0
Bit	Description										
[3]	Link 3										
[2]	Link 2										
[1]	Link 1										
[0]	Link 0										
18	DramParEn. Reset: 0.										
17	HtDataEn. Reset: 0.										
16	ProtEn. Reset: 0.										
15	L3ArrayUCEn. Reset: 0.										
14	L3ArrayCorEn. Reset: 0.										
13	Reserved.										
12	WDTRptEn. Reset: 0.										
11	AtomicRMWEn. Reset: 0.										
10	GartTblWkEn. Reset: 0. BIOS: 1. BIOS is recommended to mask logging GART table walk errors.										
9	TgtAbortEn. Reset: 0.										
8	MstrAbortEn. Reset: 0.										
7:5	SyncPktEn[2:0]. Reset: 0. SyncPktEn[3:0] = {SyncPktEn[3], SyncPktEn[2:0]}. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[3]</td> <td>Link 3</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[0]</td> <td>Link 0</td> </tr> </tbody> </table>	Bit	Description	[3]	Link 3	[2]	Link 2	[1]	Link 1	[0]	Link 0
Bit	Description										
[3]	Link 3										
[2]	Link 2										
[1]	Link 1										
[0]	Link 0										
4:2	CrcErrEn[2:0]. Reset: 0. CrcErrEn[3:0] = {CrcErrEn[3], CrcErrEn[2:0]}. <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[3]</td> <td>Link 3</td> </tr> <tr> <td>[2]</td> <td>Link 2</td> </tr> <tr> <td>[1]</td> <td>Link 1</td> </tr> <tr> <td>[0]</td> <td>Link 0</td> </tr> </tbody> </table>	Bit	Description	[3]	Link 3	[2]	Link 2	[1]	Link 1	[0]	Link 0
Bit	Description										
[3]	Link 3										
[2]	Link 2										
[1]	Link 1										
[0]	Link 0										
1	UECCEn. Reset: 0.										
0	CECCEn. Reset: 0.										

MSRC001_0049 EX Machine Check Control Mask (MC5_CTL_MASK)

Reset: 0000_0000_0000_0000h. BIOS: 0000_0000_0000_0000h. See 2.13.1 [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 2.13.1 [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 297: [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. 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. Reset: 0. 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. Reset: 0.
6	Reserved.
5	SmiEn2: SMI enable for the trap specified by MSRC001_0052. Read-write. Reset: 0.
4	Reserved.
3	SmiEn1: SMI enable for the trap specified by MSRC001_0051. Read-write. Reset: 0.
2	Reserved.
1	SmiEn0: SMI enable for the trap specified by MSRC001_0050. Read-write. Reset: 0.
0	Reserved.

MSRC001_0055 Interrupt Pending

[SharedNC](#). This register is used to specify messages that the processor generates under certain conditions, that target the IO hub. One purpose is to ensure that the IO hub can wake the processor out of the stop-grant state when there is a pending interrupt. Otherwise, it is possible for the processor to remain in the stop-grant state while an interrupt is pending in the processor. This is accomplished by sending a message to the IO hub to indicate that the interrupt is pending. There are two message types: a programmable IO-space message and the link INT_PENDING message defined by the link specification.

If the IO hub does not support the INT_PENDING message, the IO space message should be selected by IntPndMsg. When this is enabled, the check for a pending interrupt is performed at the end of each IO instruction. If there is a pending interrupt and STPCLK is asserted, the processor executes a byte-size IO access as specified by IORd, IOMsgAddr, and IOMsgData.

If the IO hub supports the INT_PENDING message, it should be selected by IntPndMsg. The check for a pending interrupt is performed while in the stop-grant state or when entering the stop-grant state. If there is a pending interrupt, the processor broadcasts the INT_PENDING message. An INT_PENDING message may not be generated for arbitrated interrupts in multi-node systems.

Bits	Description
63:32	RAZ.
31	Reserved.
30	Reserved.
29	BmStsClrOnHaltEn: BM_STS clear on Halt enable. Read-write. Reset: 0. BIOS: 1. 1=The core clears BM_STS in the IO hub by writing 0010h to the address specified by IOMsgAddr after the caches are flushed in a non-C0 C-state. This bit has no effect if cache flushing is not enabled in the C-state. See D18F3xDC[CacheFlushOnHaltCtl] and D18F4x11[C:8][CacheFlushEnCstAct0, CacheFlushEnCstAct1, CacheFlushEnCstAct2] . BmStsClrOnHaltEn must be 0 if D18F3xD4[MTC1eEn] is 0. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E] .

28:27	Reserved.
26	IORd: IO Read. Read-write. Reset: 0. 1=IO read. 0=IO write.
25	IntPndMsg: interrupt pending message. Read-write. Reset: 0. Selects the interrupt pending message type. 0=Link-defined INT_PENDING message; 1=Programmable SMI-trigger IO-space message. The status is stored in SMMFEC4[IntPendSmiSts] .
24	IntPndMsgDis: interrupt pending message disable. Read-write. Reset: 0. Disable generating the interrupt pending message specified by IntPndMsg.
23:16	IOMsgData: IO message data. Read-write. Reset: 0. IO write message data. This field is only used if IORd specifies an IO write message.
15:0	IOMsgAddr: IO message address. Read-write. Reset: 0. IO space message address. See 2.5.3.5.2 [BIOS Requirements to Initialize Message Triggered C1E] .

MSRC001_0056 SMI Trigger IO Cycle

Not-same-for-all. 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 [2.8 \[Configuration Space\]](#) 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.
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.			
	<u>Bits</u>	<u>Description</u>	<u>Bits</u>	<u>Description</u>
	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_0060 BIST Results

Read; GP-write. Reset: 0000_0000_xxxx_xxxxh. This register provides BIST results after each reset. The results provided here are identical to the values provided in EAX. If ([31:0]==0000_0000h) then no BIST failures were detected.

Bits	Description
63:0	Reserved.

MSRC001_0061 P-state Current Limit

Read; GP-write; [SharedC](#); updated-by-hardware. See [2.5.2 \[P-states\]](#).

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 2.5.2.1.2.1 [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 2.5.2.1.2.1 [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 [P-states] and 2.5.2.1.2.1 [Software P-state Numbering] .

MSRC001_0063 P-state Status

Read; GP-write; **SharedC**; **Updated-by-hardware**.

Bits	Description
63:3	RAZ.
2:0	CurPstate: current P-state. Cold reset values vary 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] and D18F3xC4 [SBI P-state Limit] ; see 2.5.2.1.7 [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 2.5.2 [P-states] and 2.5.2.1.2.1 [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 298: [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 [2.5.2 \[P-states\]](#).

When [D18F4x15C\[BoostLock\]=1](#), [MSRC001_00\[6B:64\]\[CpuVid, CpuDid, CpuFid\]](#) have special write

requirements associated with them.

Table 299: P-state Definitions

Term	Definition
CoreCOF	Core current operating frequency in MHz. CoreCOF = 100 * (MSRC001_00[6B:64][CpuFid] + 10h) / (2^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 field 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 _PSS objects (see 2.5.2.1.11 [ACPI Processor P-state Objects]) and to perform the 2.5.2.1.9 [Processor-Systemboard Power Delivery Compatibility Check]. The values are expressed in amps; they are not intended to convey final product power levels; they may not match the power levels specified in the Power and Thermal Datasheets. 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="267 976 917 1144"> <thead> <tr> <th>IddDiv</th> <th>Description</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>	IddDiv	Description	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
IddDiv	Description										
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 2.5.2.2 [NB P-states] and D18F5x170[NbPstateThreshold, NbPstateLo, NbPstateHi].										
21:16	RAZ.										
15:9	CpuVid: core VID. Read-write; Not-same-for-all. See 2.5.1 [Processor Power Planes And Voltage Control]. Writes outside the MSRC001_0071[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. This field may be modified for a given P-state when no core is in that given P-state.										

8:6	<p>CpuDid: core divisor ID. Read-write. Specifies the core frequency divisor; see CpuFid. Writes of reserved values are ignored.</p> <table border="1"> <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	<p>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. CpuFid must be less than or equal to 2Fh.</p>														

MSRC001_0070 COFVID Control

Updated-by-hardware. 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 2.5.2 [P-states].

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].
21:19	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 2.5.2.1.2.2 [Hardware P-state Numbering] .
15:9	CpuVid: core VID. Read-write. See MSRC001_00[6B:64][CpuVid].
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

See 2.5.2 [P-states].

Bits	Description
63:59	MaxNbCof: maximum NB COF. Read-only; Per-node. Reset: Product-specific. Specifies the maximum NB COF supported by the processor. The maximum frequency is 200 MHz * MaxNbCof, if MaxNbCof is greater than zero; 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. Read-only; Per-node; Updated-by-hardware. Value: Product-specific. Provides the current lowest-performance P-state limit number. This register uses hardware P-state numbering. See MSRC001_0061[CurPstateLimit] and 2.5.2.1.2.2 [Hardware P-state Numbering].
55	Reserved.
54:49	MaxCpuCof: maximum core COF. Read-only; Per-node. 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.
48:42	MinVid: minimum voltage. Read-only; Per-node. 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].
41:35	MaxVid: maximum voltage. Read-only; Per-node. 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].
34:32	StartupPstate: startup P-state number. Read-only; Per-node. 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 2.5.2.1.2.2 [Hardware P-state Numbering].
31:25	CurNbVid: current NB VID. Read-only; Updated-by-hardware. Cold reset: Product-specific. Specifies the current VDDNB voltage. See D18F5x174[CurNbVid].
24	Reserved.
23	NbPstateDis: NB P-states disabled. Read-only. Value: D18F5x174[NbPstateDis]. See 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.
22:19	Reserved.
18:16	CurPstate: current P-state. Read-only; Not-same-for-all; 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 2.5.2.1.2.2 [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: current core VID. Read-only; Updated-by-hardware. Cold Reset: Product-specific. Specifies the current VDD voltage.

8:6	CurCpuDid: current core divisor ID. Read-only; Not-same-for-all ; 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. Read-only; Not-same-for-all ; 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_0072 SBI P-state Limit

SBI-only; **Per-node**. See 2.13.4 [Sideband Interface (SBI)]. If (D18F3xE8[MultiNodeCpu]==1) then D18F3xC4 is only programmed on internal node 0 (D18F3xE8[IntNodeNum]==00b); D18F3xC4 on internal node 1 should not be written. This register specifies a P-state limit for all cores in the processor. See 2.5.2 [P-states].

Bits	Description
63:19	Reserved.
18:16	NumBoostStates: number of boosted states. Read-only. Value: D18F4x15C[NumBoostStates]. Specifies the number of P-states that are considered boosted P-states. See 2.5.2.1.1 [Application Power Management (APM)]. Provides APM read the ability to convert between hardware and software P-states for MSRC001_0072[PstateLimit] and MSRC001_003E[HtcPstateLimit] .
15:11	Reserved.
10:8	PstateLimit: P-state limit select. Read-write from APM interface. Reset: 0. Specifies a hardware P-state limit for all cores when activated by PstateLimitEn. Not changed on a write if the value written is greater than MSRC001_0061[PstateMaxVal] . No P-state limit is applied if the value written is less than D18F4x15C[NumBoostStates] . See D18F3xC4[PstateLimit] .
7:1	Reserved.
0	PstateLimitEn: P-state limit enable. Read-write from APM interface. Reset: 0. 1=PstateLimit is enabled. See D18F3xC4[PstateLimitEn] .

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 [C-state Request Interface], and D18F4x11[C:8] .

MSRC001_0075 APLM TDP Limit

SBI-only; Per-node.

Bits	Description										
63:22	Unused.										
21:16	<p>ApmITdpLimitPercent. Read-write. Reset: 0. Specifies the percentage of processor TDP allowed by APLM. See ApmITdpLimit and 2.5.2.1.1.1 [TDP Limiting].</p> <ul style="list-style-type: none"> E.g. Where the desired processor power limit in watts is ApmITdpLimitWatts[23:0], using the same bits of significance described in MSRC001_0077 (8 bits to the left of the decimal point and 16 bits to the right of the decimal point), then ApmITdpLimitPercent can be calculated as follows: <ul style="list-style-type: none"> ApmITdpLimitPercent = $\text{FLOOR}((\text{ApmITdpLimitWatts}[23:0]/(\text{MSRC001_0077}[\text{ProcessorTdp}] * \text{MSRC001_0077}[\text{Tdp2Watt}[15:0]])) * 63)$. Note that the result should be bounded between 01h and 3Fh and that 00h has special meaning. ApmITdpLimitPercent should only be accessed on internal node 0 (D18F3xE8[IntNodeNum]=00b) if D18F3xE8[MultiNodeCpu]=1. The ApmITdpLimitPercent must be greater than $(\text{MSRC001_0077}[\text{BaseTdp}]/\text{MSRC001_0077}[\text{ProcessorTdp}]) * 64$. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Disabled or 100%.</td> </tr> <tr> <td>01h</td> <td>1.6 % or $(1/64 * 100)$ %</td> </tr> <tr> <td>3Eh-02h</td> <td>$(\langle \text{ApmITdpLimitPercent} \rangle / 64) * 100$ %</td> </tr> <tr> <td>3Fh</td> <td>98.4 % or $(63/64 * 100)$ %</td> </tr> </tbody> </table>	Bits	Description	00h	Disabled or 100%.	01h	1.6 % or $(1/64 * 100)$ %	3Eh-02h	$(\langle \text{ApmITdpLimitPercent} \rangle / 64) * 100$ %	3Fh	98.4 % or $(63/64 * 100)$ %
Bits	Description										
00h	Disabled or 100%.										
01h	1.6 % or $(1/64 * 100)$ %										
3Eh-02h	$(\langle \text{ApmITdpLimitPercent} \rangle / 64) * 100$ %										
3Fh	98.4 % or $(63/64 * 100)$ %										
15:12	Unused.										
11:0	<p>ApmITdpLimit. Read-only. Reset: 0. Specifies the maximum sum of TDP for all cores on the processor allowed by APLM. If the consumed power exceeds the ApmITdpLimit, 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. See 2.5.2.1.1.1 [TDP Limiting]. A value of 0 is returned if ApmITdpLimit is disabled.</p>										

MSRC001_0077 Processor Power in TDP

Read-only; SBI-only; Per-node. BasePwrWatts, ProcessorPwrWatts, and CurrPwrWatts are fixed point integers with 8 bits to the left of the decimal point and 16 bits to the right of the decimal point; The value in W is defined by the following equation: $(\text{*PwrWatts}[23:0] / (2^{16})) \text{ W}$.

Power in watts can be calculated as follows:

- BasePwrWatts[23:0]:
 - Specifies in watts the maximum amount of power consumed by the processor for NB and logic external to the core.
 - $\text{BasePwrWatts}[23:0] = (\text{BaseTdp}[15:0] * \text{Tdp2Watt}[15:0])$. MSB's above [23] of multiply result must be 0 because result is not allowed to be $\geq 256 \text{ W}$.
- ProcessorPwrWatts[23:0]:
 - Specifies in watts the maximum amount of power the processor can support.
 - $\text{ProcessorPwrWatts}[23:0] = (\text{ProcessorTdp}[15:0] * \text{Tdp2Watt}[15:0])$. MSB's above [23] of multiply result must be 0 because result is not allowed to be $\geq 256 \text{ W}$.
- CurrPwrWatts[23:0]:
 - Specifies in watts the current amount of power being consumed by the processor.
 - CurrPwrWatts[23:0] calculation:
 - $\text{TempTdp} = \text{MSRC001_0078}[\text{ApmTdpLimit}] - (\text{signed})(\text{MSRC001_0078}[\text{TdpRunAvgAccCap}] / (2^{(\text{MSRC001_0078}[\text{RunAvgRange}]+1)))) + \text{BaseTdp}$.

- $\text{CurrPwrWatts}[23:0] = (\text{TempTdp} * \text{Tdp2Watt}[15:0])$.
- If **MSRC001_0078**[TdpRunAvgAccCap] is saturated at the largest positive or negative value, then the reported TDP is inaccurate and **D18F5xE0**[RunAvgRange] should be decreased to a smaller value.
- If **D18F3xE8**[MultiNodeCpu]=1, this field must be calculated only from internal node 0 (**D18F3xE8**[IntNodeNum]=00b).

Bits	Description
63:48	Reserved.
47:42	Tdp2Watt[5:0] . Value: D18F5xE8 [Tdp2Watt[5:0]].
41:32	Tdp2Watt[15:6] . Value: D18F5xE8 [Tdp2Watt[15:6]].
31:16	BaseTdp . Value: D18F4x1B8 [BaseTdp].
15:0	ProcessorTdp . Value: D18F4x1B8 [ProcessorTdp].

MSRC001_0078 Power Averaging Period

Read-only; SBI-only, Per-node. See [MSRC001_0077 \[Processor Power in TDP\]](#).

Bits	Description
63:45	Reserved.
44:32	ApmTdpLimit . Value: D18F5xE8 [ApmTdpLimit].
31:26	Reserved.
25:4	TdpRunAvgAccCap: processor TDP running average accumulator capture . Value: D18F5xE0 [TdpRunAvgAccCap].
3:0	RunAvgRange: running average range . Value: D18F5xE0 [RunAvgRange].

MSRC001_0079 DRAM Controller Command Throttle

SBI-only; Per-node. Reset: 0. See [2.10.10 \[DRAM On DIMM Thermal Management and Power Capping\]](#).

Bits	Description
63:24	Unused.
23:20	BwCapCmdThrottleMode: bandwidth capping command throttle mode . See D18F2xA4 [BwCapCmdThrottleMode].
19:12	Unused.
11	BwCapEn: bandwidth capping enable . See D18F2xA4 [BwCapEn].
10:0	Unused.

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 han-

pler. 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 control as follows:

- If [A, T]Valid=0, then the address range is accessed as specified by MTRRs, regardless of whether the CPU is in SMM or not.
- 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.

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 265 [Valid Values for Memory Type].
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 265 [Valid Values for Memory Type].
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)

Bits	Description
63:32	Reserved.
31:5	MBZ.

4	SvmeDisable: SVM disable. See Lock for the access type of this field. Reset: 0. 1= MSRC000_0080 [SVME] is MBZ. 0= MSRC000_0080 [SVME] is read-write. Attempting to set this field when (MSRC000_0080 [SVME]==1) causes a #GP fault, regardless of the state of Lock. See the APM2 section titled “Enabling SVM” for software use of this field.
3	Lock: SVM lock. Read; write-1-only; cleared-by-hardware. Reset: 0. See MSRC001_0118 [SvmLockKey] for the condition that causes hardware to clear this field. 1= SvmeDisable is read-only. 0= SvmeDisable is read-write.
2	DisA20m: disable A20 masking. Read-write; set-by-hardware. Reset: 0. 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. Reset: 0. 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. Reset: 0. 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)

Bits	Description
63:48	MBZ.
47:12	VM_HSAVE_PA: physical address of host save area. Read-write. Reset: 0. 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 modify SvmLockKey. If ((MSRC001_0114[Lock] ==1) && (SvmLockKey!=0) && (The write value==The value stored in SvmLockKey)) for a write to this register then hardware updates MSRC001_0114[Lock] =0.

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 revision guide for the definition of this field. See 1.2 [Reference Documents] .

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 revision guide for the definition of this field. See 1.2 [Reference Documents] .

MSRC001_020[A,8,6,4,2,0] Performance Event Select (PERF_CTL[5:0])

Reset: 0000_0000_0000_0000h. See 2.7.1 [Core Performance Monitor Counters]. See MSRC001_00[03:00]. 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 300: 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 [LVT Performance Monitor] when the performance counter overflows.										

19	Reserved.										
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 3.15 [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])

See MSRC001_020[A,8,6,4,2,0] [Performance Event Select (PERF_CTL[5:0])]. See MSRC001_00[07:04].

Table 301: 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. Reset: 0.

MSRC001_024[6,4,2,0] Northbridge Performance Event Select (NB_PERF_CTL[3:0])

Per-node. See D18F5x[70,60,50,40] and D18F5x[74,64,54,44].

Table 302: Register Mapping for MSRC001_024[6,4,2,0]

Register	Function
MSRC001_0240	Counter 0

Table 302: Register Mapping for MSRC001_024[6,4,2,0]

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. [3.16 \[NB Performance Counter Events\]](#) 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, crossbar events, HyperTransport™ interface events, and L3 cache events as documented in [3.16 \[NB Performance Counter Events\]](#). Since the northbridge performance counter register are shared by all cores on a node, monitoring of northbridge events should only be performed by one core on a node.

Bits	Description
63:41	Reserved.
40:37	IntCoreSel: interrupt to core select. Read-write. Reset: 0h. 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. Reset: 0. 1=Interrupt to a single core specified by IntCoreSel. 0=Interrupt to all cores.
35:32	EventSelect[11:8]: performance event select. Read-write. Reset: 0. See EventSelect[7:0].
31:23	Reserved.
22	En: enable performance counter. Read-write. Reset: 0. 1= Performance event counter is enabled.
21	Reserved.
20	Int: enable APIC interrupt. Read-write. Reset: 0. 1=APIC performance counter LVT interrupt is enabled to generate an interrupt via APIC340 [LVT Performance Monitor] to all APIC's on this node when the performance counter overflows.
19:16	Reserved.

15:8	UnitMask: event qualification. Read-write. Reset: 0. 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. Reset: 0. 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 3.16 [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])

Per-node. See D18F5x[78,68,58,48] and D18F5x[7C,6C,5C,4C].

Table 303: 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. Reset: 0. Returns the current value of the event counter.

3.14 MSRs - MSRC001_1xxx

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	Reserved.
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_EDX[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.BIOS: 1.
55	PerfCtrExtCore.BIOS: 1.
54	TopologyExtensions.
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	RDTSCP.
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_1020 Load-Store Configuration (LS_CFG)

Bits	Description
63:29	Reserved.
28	DisSS. Read-write; Same-for-all. 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.

Bits	Description
63:40	Reserved.
39	IF (PROC >= OR_C0) THEN DisLoopPredictor. Read-write. Reset: 0. BIOS: 1. 1=Disable loop predictor. ELSE Reserved. ENDIF.
39	Reserved.
38:10	Reserved.

9	DisSpecTlbRld. Read-write. Reset: 0. 1=Disable speculative IC TLB reload request; the request is not made to the TLB walker until the fetch is non-speculative. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot] .								
8:5	Reserved.								
4:1	<p>IF (PROC>=OR_C0) THEN</p> <p>DisIcWayFilter: disable IC way access filter. Read-write. Reset: 0. BIOS: Fh.</p> <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> <p>ELSE</p> <p>Reserved. Read-write. Controls no hardware.</p> <p>ENDIF.</p>	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
63:16	Reserved.
15	DisPfHwForSw. Read-write. Reset: 0. 1=Disable hardware prefetches for software prefetches.
14	Reserved.
13	DisHwPf. Read-write. Reset: 0. 1=Disable the DC hardware prefetcher. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot] .
12:10	Reserved.
9:5	Reserved.
4	DisSpecTlbRld. Read-write. Reset: 0. 1=Disable speculative TLB reloads. BIOS: See 2.3.3 [Using L2 Cache as General Storage During Boot] .
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 thus evictions, for L2 ways >=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	<p>L2FirstLockedWay: first L2 way locked. Read-write. Reset: 0h. See L2WayLock.</p> <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 <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 <L2FirstLockedWay> to 15 locked.	Fh	Way 15 locked.
Bits	Description								
0h	Reserved								
Eh-1h	Ways <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	IF (PROC>=OR_C0) THEN DiDtCfg4 . Read-write. Reset: 111b. BIOS: D18F3x1FC [DiDtCfg4]. ELSE Reserved. ENDIF.
41	IF (PROC>=OR_C0) THEN DiDtCfg5 . Read-write. Reset: 0. BIOS: D18F3x1FC [DiDtCfg5]. ELSE Reserved. ENDIF.
40	IF (PROC<OR_C0) THEN DiDtCfg3 . Read-write. Reset: 0. BIOS: D18F3x1FC [DiDtCfg3]. ELSE Reserved. ENDIF.
44:40	Reserved.
39:35	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].
15:8	Reserved.
7:0	Reserved.

MSRC001_1029 Decode Configuration (DE_CFG)

SharedC.

Bits	Description														
63:11	Reserved.														
10	ResyncPredSingleDispDis. Read-write. Reset: 0. BIOS: IF (PROC>=OR_C0) THEN 0 ELSE 1 ENDIF1. 1=Disables the single dispatch of loads that hit in the resync predictor.														
9	Reserved.														
8	Reserved.														
7:2	ConvertPrefetchToNop. Read-write. Reset: 0. 1=The specified PREFETCH instruction is converted into a NOP. Each of these bits apply to a different instruction as follows: <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>[0]</td> <td>PREFETCH NTA 0F 18 /0</td> </tr> <tr> <td>[1]</td> <td>PREFETCH T0 0F 18 /1</td> </tr> <tr> <td>[2]</td> <td>PREFETCH T1 0F 18 /2</td> </tr> <tr> <td>[3]</td> <td>PREFETCH T2 0F 18 /3</td> </tr> <tr> <td>[4]</td> <td>PREFETCHW mem8 0F 0D {/1, /3}</td> </tr> <tr> <td>[5]</td> <td>PREFETCH mem8 0F 0D {/0, /2, /4, /5, /6, /7}</td> </tr> </tbody> </table>	Bit	Description	[0]	PREFETCH NTA 0F 18 /0	[1]	PREFETCH T0 0F 18 /1	[2]	PREFETCH T1 0F 18 /2	[3]	PREFETCH T2 0F 18 /3	[4]	PREFETCHW mem8 0F 0D {/1, /3}	[5]	PREFETCH mem8 0F 0D {/0, /2, /4, /5, /6, /7}
Bit	Description														
[0]	PREFETCH NTA 0F 18 /0														
[1]	PREFETCH T0 0F 18 /1														
[2]	PREFETCH T1 0F 18 /2														
[3]	PREFETCH T2 0F 18 /3														
[4]	PREFETCHW mem8 0F 0D {/1, /3}														
[5]	PREFETCH mem8 0F 0D {/0, /2, /4, /5, /6, /7}														
1:0	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.
49:43	Reserved.
42	ProbeFilterSupEn. Read-write. Reset: 0. BIOS: See 2.9.4.1 [Probe Filter] . 1=Enable probe filter support. Read probes transition E to O. Move data for E lines. E state lines can't silently transition to S or be dropped.
41:38	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. Required to be set for the Monitor/MWait instructions.
9:8	Reserved.

7:6	<p>ThrottleNbInterface[1:0]. Read-write. ThrottleNbInterface[3:0] = {ThrottleNbInterface[3:2], ThrottleNbInterface[1:0]}. Reset: 11b. BIOS: NumOfCompUnitsOnNode-1. 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.</p> <p>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].</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>0 Clocks.</td> </tr> <tr> <td>1h</td> <td>1 Clock.</td> </tr> <tr> <td>2h</td> <td>2 Clocks.</td> </tr> <tr> <td>3h</td> <td>3 Clocks.</td> </tr> <tr> <td>Fh-4h</td> <td>Reserved.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	0h	0 Clocks.	1h	1 Clock.	2h	2 Clocks.	3h	3 Clocks.	Fh-4h	Reserved.
<u>Bits</u>	<u>Description</u>												
0h	0 Clocks.												
1h	1 Clock.												
2h	2 Clocks.												
3h	3 Clocks.												
Fh-4h	Reserved.												
5:0	Reserved.												

MSRC001_102B Combined Unit Configuration 3 (CU_CFG3)

SharedC.

Bits	Description										
63:53	Reserved.										
52:51	<p>AsidDecrScaleFactor. Read-write. Reset: 00b. Specifies the decrement rate for the ASID replacement counter for ASIDs not currently in use; The larger the value programmed the slower the counter decrements.</p> <table border="1"> <thead> <tr> <th><u>Bits</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Inactive ASID replacement counter decrements every 16 TLB inserts.</td> </tr> <tr> <td>01b</td> <td>Inactive ASID replacement counter decrements every 32 TLB inserts.</td> </tr> <tr> <td>10b</td> <td>Inactive ASID replacement counter decrements every 64 TLB inserts.</td> </tr> <tr> <td>11b</td> <td>Inactive ASID replacement counter decrements every 128 TLB inserts.</td> </tr> </tbody> </table>	<u>Bits</u>	<u>Description</u>	00b	Inactive ASID replacement counter decrements every 16 TLB inserts.	01b	Inactive ASID replacement counter decrements every 32 TLB inserts.	10b	Inactive ASID replacement counter decrements every 64 TLB inserts.	11b	Inactive ASID replacement counter decrements every 128 TLB inserts.
<u>Bits</u>	<u>Description</u>										
00b	Inactive ASID replacement counter decrements every 16 TLB inserts.										
01b	Inactive ASID replacement counter decrements every 32 TLB inserts.										
10b	Inactive ASID replacement counter decrements every 64 TLB inserts.										
11b	Inactive ASID replacement counter decrements every 128 TLB inserts.										
50	<p>AsidIncrScaleFactor. Read-write. Reset: 0. Specifies the increment rate for the ASID replacement counter for ASIDs currently in use; The larger the value programmed the slower the counter increments. 0=Active ASID replacement counter increments every 16 TLB inserts. 1=Active ASID replacement counter increments every 64 TLB inserts.</p>										
49	<p>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 during boot; See 2.3.3 [Using L2 Cache as General Storage During Boot]; Must either be set according to 2.9.4.2 [Probe Filter and ATM Mode Initialization Sequence] or be set before passing control to the OS. 1=The effective host 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; does not affect guest CR0[CD]; SMM taken from guest mode must coordinate with the other core of the compute unit if the MTRR's are changed.</p>										
48:43	Reserved.										
42	<p>PwcDisableWalkerSharing. Read-write. Reset: 0. BIOS: IF (PROC>=OR_C0) THEN 0 ELSE 1 ENDIF1. 1=Page table walker sharing is disabled. Core 0 uses page walker 0 and Core 1 uses page walker 1.</p>										
41:23	Reserved.										

22	IF (PROC>=OR_C0) THEN 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. ELSE Reserved. ENDIF.										
21:20	IF (PROC>=OR_C0) THEN 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> ELSE Reserved. ENDIF.	Bits	Description	00b	3	01b	4	10b	5	11b	6
Bits	Description										
00b	3										
01b	4										
10b	5										
11b	6										
22:20	Reserved.										
19	Reserved.										
18	PfcDis. Read-write. Reset: 0. 1=Prefetcher disabled.										
17	PfcStrideDis. Read-write. Reset: 0. 1=Stride prefetch generation disabled.										
16	PfcRegionDis. Read-write. Reset: 0. 1=Region prefetch generation disabled.										
15:4	Reserved.										
3	PfcL1TrainDis: stride training to L1 disable. Read-write. Reset: 0. 1=L1 prefetch training disabled.										
2:0	Reserved.										

MSRC001_102C Execution Unit Configuration (EX_CFG)

Bits	Description
63:55	Reserved.
54	IF (PROC>=OR_C0) THEN LateSbzResync. Read-write. Reset: 0. BIOS: 1. 1=Force an SBZ (Shift-By-Zero) resync to wait if the op has a store component which has not yet completed. 0=The SBZ resync is taken as soon as EX detects the condition, without waiting for any other component to complete. ELSE Reserved. ENDIF.
54	Reserved.
53:0	Reserved.

IF (PROC>=OR_C0) THEN

MSRC001_102D Load-Store Configuration 2 (LS_CFG2)

Reset: 0000_0000_0202_0000h.

Bits	Description
63:24	Reserved.
23	DisScbThreshold. Read-write. Reset: 0. BIOS: 1. 1=Disable SCB threshold and force deallocation.
22:15	Reserved.
14	ForceSmcCheckFlwStDis. Read-write. Reset: 0. BIOS: 1. 0=Force a self modifying code check when a cache probe hits a store that has not retired. 1=Do not perform a self modifying code check when a cache probe hits a store that has not retired.
13	Reserved.
12	ForceBusLockDis. Read-write. Reset: 0. BIOS: 1. 1=Disable widget that forces cacheable locks to be bus locks.
11:0	Reserved.

ENDIF.

MSRC001_1030 IBS Fetch Control (IC_IBS_CTL)

Reset: 0000_0000_0000_0000h. See 2.7.3 [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. 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 2.7.3 [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. Read-write. 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; IbsOpCurCnt[26:0] is cleared when IbsOpEn is changed from 0 to 1.						
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="0" style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: left;"><u>Bits</u></td> <td style="text-align: left;"><u>Description</u></td> </tr> <tr> <td>0008h-0000h</td> <td>Reserved</td> </tr> <tr> <td>FFFFh-0009h</td> <td><IbsOpMaxCnt[19:4]*16> ops</td> </tr> </table>	<u>Bits</u>	<u>Description</u>	0008h-0000h	Reserved	FFFFh-0009h	<IbsOpMaxCnt[19:4]*16> ops
<u>Bits</u>	<u>Description</u>						
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 retired. 1=Tagged operation was a branch 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	IbsOpMispReturn: mispredicted return micro-op. 1=Tagged operation was a mispredicted return micro-op. Qualified by IbsOpBrnRet==1.
32	IbsOpBrnResync: resync micro-op. 1=Tagged operation was resync micro-op.
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
63:32	Reserved.
31:6	Reserved.
5	NbIbsReqCacheHitSt: IBS cache hit 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.
	<u>Bits</u> <u>Description</u>
0h	No valid status
1h	L3: data returned from local L3 cache
2h	Cache: data returned from another compute-unit cache on the same node or an L3 or compute unit cache on a different 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_0001_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. BIOS: 0. 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	IbsDcStToLdCan: data forwarding from store to load operation canceled. 1=Data forwarding from a store operation to the tagged load was canceled.
11	IbsDcStToLdFwd: data forwarded from store to load operation. 1=Data for tagged load operation was forwarded from a store operation. If this bit is set and IbsDcStToLdCan=1, then the data for the load operation forwarded from a store operation but the data was not forwarded immediately.
10	Reserved.
9	IbsDcLdBnkCon: bank conflict on load operation. 1=A bank conflict with a load operation occurred in the data cache on the tagged load or store operation.

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)

Bits	Description
63:48	RAZ.
47:0	IbsDcPhysAd: load or store physical address. Read-write; updated-by-hardware. Reset: 0. 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

Bits	Description
63:9	Reserved.
8	LvtOffsetVal: local vector table offset valid. Read-only; GP-write ; Per-node . Reset: 0. See D18F3x1CC [LvtOffsetVal].
7:4	Reserved.
3:0	LvtOffset: local vector table offset. Read-only; GP-write ; Per-node . Reset: 0. See D18F3x1CC [LvtOffset].

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 .

3.15 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 [2.7.2 \[NB Performance Monitor Counters\]](#) and [3.16 \[NB Performance Counter Events\]](#).

3.15.1 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 Floating Point Ops

[PERF_CTL\[3\]](#). This is a retire-based event. The number of retired 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.

3.15.2 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 non-speculative execution time.

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.

IF (PROC>=OR_C0) THEN

PMCx032 Misaligned Stores

PERF_CTL[5:0]. Implemented by LS. The number of misaligned stores.

ENDIF.

IF (PROC>=OR_C0) THEN

PMCx034 FP +Load Buffer Stall

PERF_CTL[5:0]. Implemented by LS. The number of loads stalled due to buffer full.

ENDIF.

3.15.3 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.

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**.

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.

3.15.4 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>Bits</th> <th>Description</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>	Bits	Description	27-0	MAB ID	255-28	Reserved
Bits	Description						
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 (if modified unwritten is disabled by D18F0x68 [ATMMoDeEn]==0); Modified unwritten (if modified unwritten is enabled by D18F0x68 [ATMMoDeEn]==1).
0	Exclusive

PMCx06D Octwords 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.

PMCx075 Cache Cross-invalidates

PERF_CTL[2:0]. These reflect internal probes for Icache or Dcache misses that hit in the Dcache or Icache, causing the line to be invalidated. These may result from code modification, or data being located too close to code, or virtual address aliasing. The aliasing cases arise when a physical memory location is referenced via two or more virtual addresses which differ in bits 14:12.

UnitMask	Description
7:4	Reserved.
3	IC Invalidates DC (execution of recently modified code, or modified data too close to code).
2	IC Invalidates IC (aliasing)
1	DC Invalidates DC (aliasing)
0	DC Invalidates IC (modification of cached instructions, or of data located too close to code).

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.

3.15.5 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 to the L2. This event is not core specific and for either core counts the IC victims caused by both cores of the compute unit.

PMCx08C Instruction Cache Lines Invalidated

PERF_CTL[2:0]. The number of instruction cache lines invalidated. A non-SMC event is CMC (cross modifying code), either from the other core of the compute unit or another compute compute unit.

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.

3.15.6 PMCx0[BF:A0] Events

There are no PMCx0[BF:A0] Events.

3.15.7 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 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 **PMCx0D3** to **PMCx0D9**, 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 microsequencer 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 microsequencer 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 304: 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.

Data matches:

- If (**PROC**<**OR_C0**) then dData matches are counted if the breakpoint is enabled and the data access becomes non-speculative, but not necessarily retired.
- If (**PROC**>=**OR_C0**) then data matches are counted regardless of whether the breakpoint is enabled and the data access becomes non-speculative, but not necessarily retired. This is GH and TN behavior.
- Load/store breakpoint matches do not incur any overhead.

Instruction matches:

- Instruction matches becomes non-speculative, but not necessarily retired. Instruction matches do not depend on the breakpoint being enabled.
- Each instruction breakpoint match incurs an overhead of about 120 cycles

IF (**PROC**>=**OR_C0**) THEN

PMCx1C0 Retired x87 Floating Point Operations

PERF_CTL[5:3]. Implemented by EX. 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

ENDIF.

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 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.

IF (PROC>=OR_C0) THEN

PMCx1DD Cycles Without Dispatch Due To Integer PRF Tokens

PERF_CTL[2:0]. Implemented by DE. The number of cycles a core has valid ops for dispatch and one of the reasons for a dispatch stall is the absence of sufficient integer PRF tokens. This event counts even when dispatch selects the other core of the compute-unit.

ENDIF.

IF (PROC>=OR_C0) THEN

PMCx1DE Cycles Without Dispatch Due to FP PRF Tokens

PERF_CTL[2:0]. Implemented by DE. The number of cycles a core has valid ops for dispatch and one of the reasons for a dispatch stall is the absence of sufficient FP PRF tokens. This event requires that the other core of the compute unit is in the Halt state. This event counts even when dispatch selects the other core of the compute-unit.

ENDIF.

3.16 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])].

3.16.1 NBPMCx0E[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

3.16.2 NBPMCx0E[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: Subcache-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

NBPMCx0EE GART Events

These events reflect GART activity, and in particular allow one to calculate the GART TLB miss ratio as GART_miss_count divided by GART_aperture_hit_count. GART aperture accesses are typically from IO devices as opposed to the processor, and generally from a 3D graphics accelerator, but can be from other devices when the GART is used as an IOMMU.

UnitMask	Description
7	GART multiple table walk in progress
6:4	Reserved.
3	GART Request hit table walk in progress
2	GART miss
1	GART aperture hit on access from IO
0	GART aperture hit on access from CPU

3.16.3 NBPMCx0F[F:0] Events (Link, Crossbar)

NBPMCx[1F9,0F8,0F7,0F6] Link Transmit Bandwidth

Table 305: Register Mapping for NBPMCx[1F9,0F8,0F7,0F6]

Register	Function
NBPMCx0F6	Link 0
NBPMCx0F7	Link 1
NBPMCx0F8	Link 2
NBPMCx1F9	Link 3

The number of DWs transmitted (or unused, in the case of NOPs) on the outgoing side of the links. The count for (UnitMask[7:0]==3Fh) is the maximum transmission rate of the link. Link utilization may be calculated by (The count for (UnitMask[7:0]==37h))/(The count for (UnitMask[7:0]==3Fh)), described as non-NOP traffic divided by total traffic. Bandwidth in terms of bytes per unit time for any one component or combination of components is calculated by multiplying the count by four and dividing by elapsed time. The Data event provides a direct indication of the flow of data around the system. Translating this link-based view into a source/target node based view requires knowledge of the system layout (i.e. which links connect to which nodes).

UnitMask	Description
7	SubLinkMask: sublink select. Sublink select if link is ungangled; reserved if gangled. 0=Sublink 0. 1=Sublink 1.
6	Reserved.
5	Per packet CRC sent
4	Address (including extensions) DW sent

3	NOP DW sent (idle)
2	Buffer release DW sent
1	Data DW sent
0	Command DW sent

3.16.4 NBPMCx1E[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

NBPMCx1EA Request Cache Status 0

The probe response type for RdBlk or RdBlkS request type. [D18F0x68\[ATMMoDeEn\]=1](#) is required for this event.

UnitMask	Description
7	Track Cache Stat for RdBlkS
6	Track Cache Stat for RdBlk
5	Directed Probe
4	Probe Miss
3	Probe Hit M
2	Probe Hit MuW or O
1	Probe Hit E
0	Probe Hit S

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

3.16.5 NBPMCx1F[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

NBPMCx1F9: See [NBPMCx\[1F9,0F8,0F7,0F6\]](#) [[Link Transmit Bandwidth](#)].

3.16.6 NBPMCx[5F:4E][F:0] Events (L3 Cache)

NBPMCx4E0 Read Request to L3 Cache

This event counts the read requests from each core to the L3 cache including canceled requests.

UnitMask	Description												
7:4	<p>CoreSel: core select. Selects a single core or all cores on the node to be tracked.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Core 0</td> </tr> <tr> <td>6h-1h</td> <td>Core <CoreSel></td> </tr> <tr> <td>7h</td> <td>Core 7</td> </tr> <tr> <td>Eh-8h</td> <td>Reserved</td> </tr> <tr> <td>Fh</td> <td>All cores</td> </tr> </tbody> </table>	Bits	Description	0h	Core 0	6h-1h	Core <CoreSel>	7h	Core 7	Eh-8h	Reserved	Fh	All cores
Bits	Description												
0h	Core 0												
6h-1h	Core <CoreSel>												
7h	Core 7												
Eh-8h	Reserved												
Fh	All cores												
3	1=Count prefetch only; 0=Count prefetch and non-prefetch.												
2	Read Block Modify												
1	Read Block Shared (Instruction cache read)												
0	Read Block Exclusive (Data cache read)												

NBPMCx4E1 L3 Cache Misses

This event counts the number of L3 cache misses for accesses from each core. The approximate number of L3 hits can be determined by subtracting this event from [NBPMCx4E0](#).

UnitMask	Description												
7:4	CoreSel: core select. Selects a single core or all cores on the node to be tracked. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Core 0</td> </tr> <tr> <td>6h-1h</td> <td>Core <CoreSel></td> </tr> <tr> <td>7h</td> <td>Core 7</td> </tr> <tr> <td>Eh-8h</td> <td>Reserved</td> </tr> <tr> <td>Fh</td> <td>All cores</td> </tr> </tbody> </table>	Bits	Description	0h	Core 0	6h-1h	Core <CoreSel>	7h	Core 7	Eh-8h	Reserved	Fh	All cores
Bits	Description												
0h	Core 0												
6h-1h	Core <CoreSel>												
7h	Core 7												
Eh-8h	Reserved												
Fh	All cores												
3	1=Count prefetch only; 0=Count prefetch and non-prefetch.												
2	Read Block Modify												
1	Read Block Shared (Instruction cache read)												
0	Read Block Exclusive (Data cache read)												

NBPMCx4E2 L3 Fills caused by L2 Evictions

This event counts the number of L3 fills caused by L2 evictions.

UnitMask	Description												
7:4	CoreSel: core select. Selects a single core or all cores on the node to be tracked. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Core 0</td> </tr> <tr> <td>6h-1h</td> <td>Core <CoreSel></td> </tr> <tr> <td>7h</td> <td>Core 7</td> </tr> <tr> <td>Eh-8h</td> <td>Reserved</td> </tr> <tr> <td>Fh</td> <td>All cores</td> </tr> </tbody> </table>	Bits	Description	0h	Core 0	6h-1h	Core <CoreSel>	7h	Core 7	Eh-8h	Reserved	Fh	All cores
Bits	Description												
0h	Core 0												
6h-1h	Core <CoreSel>												
7h	Core 7												
Eh-8h	Reserved												
Fh	All cores												
3	Modified												
2	Owned												
1	Exclusive												
0	Shared												

NBPMCx4E3 L3 Evictions

This event counts the state of the L3 lines when they are evicted from the L3 cache.

UnitMask	Description
7:4	Reserved.
3	Modified
2	Owned
1	Exclusive
0	Shared

NBPMCx4ED Non-canceled L3 Read Requests

This event tracks all read requests from each core to the L3 cache that are not canceled.

UnitMask	Description												
7:4	CoreSel: core select. Selects a single core or all cores on the node to be tracked. <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Core 0</td> </tr> <tr> <td>6h-1h</td> <td>Core <CoreSel></td> </tr> <tr> <td>7h</td> <td>Core 7</td> </tr> <tr> <td>Eh-8h</td> <td>Reserved</td> </tr> <tr> <td>Fh</td> <td>All cores</td> </tr> </tbody> </table>	Bits	Description	0h	Core 0	6h-1h	Core <CoreSel>	7h	Core 7	Eh-8h	Reserved	Fh	All cores
Bits	Description												
0h	Core 0												
6h-1h	Core <CoreSel>												
7h	Core 7												
Eh-8h	Reserved												
Fh	All cores												
3	1=Count prefetch only; 0=Count prefetch and non-prefetch.												
2	RdBlkM												
1	RdBlkS												
0	RdBlk												

NBPMCx4EF L3 Latency

This event enables the average latency for L3 requests to be calculated ($L3CycCount/L3ReqCount$).

UnitMask	Description
7:2	Reserved.
1	L3ReqCount. L3 request count.
0	L3CycCount. L3 Request cycle count.

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.

55	SMMFEC0: SMM IO Trap Offset	185	F2RC2: DDR3 Buffer Function 2 Control Word 2
55	SMMFEC4: Local SMI Status	185	F2RC3: DDR3 Buffer Function 2 Control Word 3
56	SMMFEC8: SMM IO Restart Byte	185	F2RC[15:8,6:4]: DDR3 Buffer Function 2 Control Word [15:8,6:4]
56	SMMFEC9: Auto Halt Restart Offset	185	F3RC0: DDR3 Buffer Function 3 Control Word 0
57	SMMFECA: NMI Mask	185	F3RC1: DDR3 Buffer Function 3 Control Word 1
57	SMMFED8: SMM SVM State	185	F3RC2: DDR3 Buffer Function 3 Control Word 2
57	SMMFEFC: SMM-Revision Identifier	186	F3RC[5:3]: DDR3 Buffer Function 3 Control Word [5:3]
58	SMMFF00: SMM Base Address (SMM_BASE)	186	F3RC6: DDR3 Buffer Function 3 Control Word 6
165	MR0_dct[1:0]: DDR3 MR0	186	F3RC8: DDR3 Buffer Function 3 Control Word 8
166	MR1_dct[1:0]: DDR3 MR1	186	F3RC9: DDR3 Buffer Function 3 Control Word 9
167	MR2_dct[1:0]: DDR3 MR2	186	F[10:3]RC10: DDR3 Buffer Function [10:3] Control Word 10
167	MR3: DDR3 MR3	186	F[10:3]RC11: DDR3 Buffer Function [10:3] Control Word 11
168	RC0: DDR3 Register Control Word 0	187	F[11:3]RC12: DDR3 Buffer Function [11:3] Control Word 12
168	RC1: DDR3 Register Control Word 1	187	F[11:3]RC13: DDR3 Buffer Function [11:3] Control Word 13
168	RC2: DDR3 Register Control Word 2	187	F3RC[15:14]: DDR3 Buffer Function 3 Control Word [15:14]
174	RC3: DDR3 Register Control Word 3	187	F[11:4]RC[15:14,9:8,6:0]: DDR3 Buffer Function [11:4] Control Word [15:14,9:8,6:0]
174	RC4: DDR3 Register Control Word 4	187	F11RC[11:10]: DDR3 Buffer Function 11 Control Word [11:10]
174	RC5: DDR3 Register Control Word 5	187	F12RC[15:8,6:0]: DDR3 Buffer Function 12 Control Word [15:8,6:0]
175	RC[7:6]: DDR3 Register Control Word [7:6]	187	F13RC[9:8,6:0]: DDR3 Buffer Function 13 Control Word [9:8,6:0]
175	RC8: DDR3 Register Control Word 8	188	F13RC10: DDR3 Buffer Function 13 Control Word 10
175	RC9: DDR3 Register Control Word 9	188	F13RC11: DDR3 Buffer Function 13 Control Word 11
175	RC10: DDR3 Register Control Word 10	188	F13RC[13:12]: DDR3 Buffer Function 13 Control Word [13:12]
175	RC11: DDR3 Register Control Word 11	188	F13RC14: DDR3 Buffer Function 13 Control Word 14
176	RC[15:12]: DDR3 Register Control Word [15:12]	188	F13RC15: DDR3 Buffer Function 13 Control Word 15
176	F0RC0: DDR3 Buffer Function 0 Control Word 0	188	F[15:14]RC[15:8,6:0]: DDR3 Buffer Function [15:14] Control Word [15:8,6:0]
176	F0RC1: DDR3 Buffer Function 0 Control Word 1	189	ExtRC_xAC: DDR3 Buffer Extended Control Word ACh
176	F0RC2: DDR3 Buffer Function 0 Control Word 2	189	ExtRC_x[BF:B8]: DDR3 Buffer Extended Control Word BFh:B8h
176	F0RC3: DDR3 Buffer Function 0 Control Word 3	189	ExtRC_xC8: DDR3 Buffer Extended Control Word C8h
177	F0RC4: DDR3 Buffer Function 0 Control Word 4	190	ExtRC_xC9: DDR3 Buffer Extended Control Word C9h
177	F0RC5: DDR3 Buffer Function 0 Control Word 5	190	ExtRC_xCA: DDR3 Buffer Extended Control Word CAh
177	F0RC6: DDR3 Buffer Function 0 Control Word 6	190	ExtRC_xCB: DDR3 Buffer Extended Control Word CBh
177	F[15:0]RC7: DDR3 Buffer Function [15:0] Control Word 7	190	ExtRC_xCC: DDR3 Buffer Extended Control Word CCh
177	F0RC8: DDR3 Buffer Function 0 Control Word 8	190	ExtRC_xCD: DDR3 Buffer Extended Control Word CDh
179	F0RC9: DDR3 Buffer Function 0 Control Word 9	191	ExtRC_xCE: DDR3 Buffer Extended Control Word CEh
179	F0RC10: DDR3 Buffer Function 0 Control Word 10	191	ExtRC_xCF: DDR3 Buffer Extended Control Word CFh
180	F0RC11: DDR3 Buffer Function 0 Control Word 11	260	IOCF8: IO-Space Configuration Address
180	F0RC12: DDR3 Buffer Function 0 Control Word 12	261	IOCF8: IO-Space Configuration Data Port
181	F0RC13: DDR3 Buffer Function 0 Control Word 13	261	D18F0x00: Device/Vendor ID
181	F0RC14: DDR3 Buffer Function 0 Control Word 14	261	D18F0x04: Status/Command
182	F0RC15: DDR3 Buffer Function 0 Control Word 15	261	D18F0x08: Class Code/Revision ID
182	F1RC0: DDR3 Buffer Function 1 Control Word 0	261	D18F0x0C: Header Type
182	F1RC1: DDR3 Buffer Function 1 Control Word 1	262	D18F0x34: Capabilities Pointer
182	F1RC2: DDR3 Buffer Function 1 Control Word 2	262	D18F0x[5C:40]: Routing Table
183	F1RC3: DDR3 Buffer Function 1 Control Word 3	263	D18F0x60: Node ID
183	F1RC[6:4]: DDR3 Buffer Function 1 Control Word [6:4]	264	D18F0x64: Unit ID
183	F1RC8: DDR3 Buffer Function 1 Control Word 8	264	D18F0x68: Link Transaction Control
183	F1RC9: DDR3 Buffer Function 1 Control Word 9	267	D18F0x6C: Link Initialization Control
183	F1RC10: DDR3 Buffer Function 1 Control Word 10	268	D18F0x[E0,C0,A0,80]: Link Capabilities
183	F1RC11: DDR3 Buffer Function 1 Control Word 11	269	D18F0x[E4,C4,A4,84]: Link Control
184	F1RC12: DDR3 Buffer Function 1 Control Word 12	271	D18F0x[E8,C8,A8,88]: Link Frequency/Revision
184	F1RC13: DDR3 Buffer Function 1 Control Word 13	273	D18F0x[EC,CC,AC,8C]: Link Feature Capability
184	F1RC14: DDR3 Buffer Function 1 Control Word 14	274	D18F0x[F0,D0,B0,90]: Link Base Channel Buffer Count
184	F1RC15: DDR3 Buffer Function 1 Control Word 15	276	D18F0x[F4,D4,B4,94]: Link Isochronous Channel Buffer Count
184	F2RC0: DDR3 Buffer Function 2 Control Word 0		
184	F2RC1: DDR3 Buffer Function 2 Control Word 1		

277	D18F0x[F8,D8,B8,98]: Link Type
277	D18F0x[FC,DC,BC,9C]: Link Frequency Extension
278	D18F0x[11C,118,114,110]: Link Clumping Enable
278	D18F0x[12C,128,124,120]: Sublink 1 Clumping Enable
279	D18F0x[14C:130]: Link Retry
280	D18F0x150: Link Global Retry Control
281	D18F0x160: Extended Node ID
282	D18F0x164: Coherent Link Traffic Distribution
283	D18F0x168: Extended Link Transaction Control
283	D18F0x16C: Link Global Extended Control
284	D18F0x[18C:170]: Link Extended Control
286	D18F0x1A0: Link Initialization Status
286	D18F0x1DC: Core Enable
287	D18F0x1E0: Coherent Link Pair Traffic Distribution
287	D18F1x00: Device/Vendor ID
288	D18F1x08: Class Code/Revision ID
288	D18F1x0C: Header Type
288	D18F1x[17C:140,7C:40]: DRAM Base/Limit
291	D18F1x[1CC:180,BC:80]: MMIO Base/Limit
294	D18F1x[DC:C0]: IO-Space Base/Limit
296	D18F1x[EC:E0]: Configuration Map
297	D18F1xF0: DRAM Hole Address
298	D18F1xF4: VGA Enable
298	D18F1x10C: DCT Configuration Select
299	D18F1x120: DRAM Base System Address
299	D18F1x124: DRAM Limit System Address
300	D18F2x00: Device/Vendor ID
300	D18F2x08: Class Code/Revision ID
300	D18F2x0C: Header Type
300	D18F2x[5C:40]_dct[1:0]: DRAM CS Base Address
302	D18F2x[6C:60]_dct[1:0]: DRAM CS Mask
303	D18F2x78_dct[1:0]: DRAM Control
304	D18F2x7C_dct[1:0]: DRAM Initialization
305	D18F2x80_dct[1:0]: DRAM Bank Address Mapping
306	D18F2x84_dct[1:0]: DRAM MRS
307	D18F2x88_dct[1:0]: DRAM Timing Low
307	D18F2x8C_dct[1:0]: DRAM Timing High
308	D18F2x90_dct[1:0]: DRAM Configuration Low
310	D18F2x94_dct[1:0]: DRAM Configuration High
312	D18F2x98_dct[1:0]: DRAM Controller Additional Data Offset
313	D18F2x9C_dct[1:0]: DRAM Controller Additional Data Port
313	D18F2x9C_x0000_0000_dct[1:0]: DRAM Output Driver Compensation Control
315	D18F2x9C_x0000_0[3:0]0[3:1]_dct[1:0]: DRAM Write Data Timing
316	D18F2x9C_x0000_0004_dct[1:0]: DRAM Address/Command Timing Control
317	D18F2x9C_x0000_0[3:0]0[7:5]_dct[1:0]: DRAM Read DQS Timing
318	D18F2x9C_x0000_0008_dct[1:0]: DRAM Phy Control
319	D18F2x9C_x0000_000B_dct[1:0]: DRAM Phy Status Register
319	D18F2x9C_x0000_000C_dct[1:0]: DRAM Phy Miscellaneous
320	D18F2x9C_x0000_000D_dct[1:0]: DRAM Phy DLL Control
321	D18F2x9C_x0000_00[2A:10]_dct[1:0]: DRAM DQS Receiver Enable Timing
322	D18F2x9C_x0000_00[4A:30]_dct[1:0]: DRAM DQS Write Timing
324	D18F2x9C_x0000_00[52:50]_dct[1:0]: DRAM Phase Recovery Control
325	D18F2x9C_x0D0F_0[F,8:0]02_dct[1:0]: Data Byte Transmit PreDriver Calibration
326	D18F2x9C_x0D0F_0[F,8:0]04_dct[1:0]: Data Byte DM Configuration
327	D18F2x9C_x0D0F_0[F,8:0]0[A,6]_dct[1:0]: Data Byte Transmit PreDriver Calibration 2
327	D18F2x9C_x0D0F_0[F,8:0]0F_dct[1:0]: Data Byte DLL Clock Enable
328	D18F2x9C_x0D0F_0[F,8:0]10_dct[1:0]: Data Byte DLL Power Management
329	D18F2x9C_x0D0F_0[F,8:0]13_dct[1:0]: Data Byte DLL Configuration
329	D18F2x9C_x0D0F_0[F,8:0]1F_dct[1:0]: Data Byte Receiver Configuration
330	D18F2x9C_x0D0F_0[F,8:0]30_dct[1:0]: Data Byte DLL Configuration and Power Down
330	D18F2x9C_x0D0F_0[F,8:0]31_dct[1:0]: Data Byte Fence2 Threshold
331	D18F2x9C_x0D0F_2[2:0]02_dct[1:0]: Clock Transmit PreDriver Calibration
332	D18F2x9C_x0D0F_[C,8,2][2:0]1F_dct[1:0]: Receiver Configuration

333 D18F2x9C_x0D0F_2[F,2:0]30_dct[1:0]: Clock DLL Configuration and Power Down
 333 D18F2x9C_x0D0F_[C,8,2][2:0]31_dct[1:0]: Fence2 Threshold
 334 D18F2x9C_x0D0F_4009_dct[1:0]: Phy Cmp Configuration
 335 D18F2x9C_x0D0F_[C,8][1:0]02_dct[1:0]: Transmit PreDriver Calibration
 336 D18F2x9C_x0D0F_[C,8][1:0][12,0E,0A,06]_dct[1:0]: Transmit PreDriver Calibration 2
 336 D18F2x9C_x0D0F_8021_dct[1:0]: DLL CS 6 & 7 Timing Control
 337 D18F2x9C_x0D0F_812F_dct[1:0]: Tristate Configuration
 337 D18F2x9C_x0D0F_E003_dct[1:0]: Phy Calibration Configuration
 338 D18F2x9C_x0D0F_E006_dct[1:0]: Phy PLL Lock Time
 338 D18F2x9C_x0D0F_E008_dct[1:0]: Phy Fence Register
 338 D18F2x9C_x0D0F_E00A_dct[1:0]: Phy Dynamic Power Mode
 338 D18F2x9C_x0D0F_E013_dct[1:0]: Phy PLL Regulator Wait Time
 338 D18F2xA0_dct[1:0]: DRAM Controller Miscellaneous
 339 D18F2xA4: DRAM Controller Temperature Throttle
 340 D18F2xA8_dct[1:0]: DRAM Controller Miscellaneous 2
 343 D18F2xAC: DRAM Controller Temperature Status
 343 D18F2xF8: P-state Power Information 1
 344 D18F2xFC: P-state Power Information 2
 344 D18F2x104: P-state Power Information 3
 344 D18F2x10C: Swap Interleaved Region Base/Limit
 345 D18F2x110: DRAM Controller Select Low
 346 D18F2x114: DRAM Controller Select High
 347 D18F2x118: Memory Controller Configuration Low
 348 D18F2x11C: Memory Controller Configuration High
 351 D18F2x1B0: Extended Memory Controller Configuration Low
 352 D18F2x1B4: Extended Memory Controller Configuration High Register
 354 D18F2x200_dct[1:0]: DRAM Timing 0
 355 D18F2x204_dct[1:0]: DRAM Timing 1
 356 D18F2x208_dct[1:0]: DRAM Timing 2
 356 D18F2x20C_dct[1:0]: DRAM Timing 3
 357 D18F2x210_dct[1:0]_nbp[3:0]: DRAM NB P-state
 358 D18F2x214_dct[1:0]: DRAM Timing 4
 359 D18F2x218_dct[1:0]: DRAM Timing 5
 360 D18F2x21C_dct[1:0]: DRAM Timing 6
 361 D18F2x220_dct[1:0]: DRAM Timing 7
 361 D18F2x224_dct[1:0]: DRAM Timing 8
 362 D18F2x228_dct[1:0]: DRAM Timing 9
 362 D18F2x22C_dct[1:0]: DRAM Timing 10
 363 D18F2x[234:230]_dct[1:0]: DRAM Read ODT Pattern
 364 D18F2x[23C:238]_dct[1:0]: DRAM Write ODT Pattern
 364 D18F2x240_dct[1:0]: DRAM ODT Control
 365 D18F2x244_dct[1:0]: DRAM Controller Miscellaneous 3
 365 D18F2x248_dct[1:0]: DRAM Power Management 0
 367 D18F2x24C_dct[1:0]: DRAM Power Management 1
 368 D18F2x250_dct[1:0]: DRAM Loopback and Training Control
 369 D18F2x25[8,4]_dct[1:0]: DRAM Target Base
 369 D18F2x260_dct[1:0]: DRAM Command 1
 370 D18F2x264_dct[1:0]: DRAM Status 0
 370 D18F2x268_dct[1:0]: DRAM Status 1
 371 D18F2x26C_dct[1:0]: DRAM Status 2
 371 D18F2x270_dct[1:0]: DRAM PRBS
 371 D18F2x274_dct[1:0]: DRAM DQ Mask Low
 372 D18F2x278_dct[1:0]: DRAM DQ Mask High
 372 D18F2x27C_dct[1:0]: DRAM ECC Mask
 372 D18F2x28C_dct[1:0]: DRAM Command 2
 373 D18F2x290_dct[1:0]: DRAM Status 3
 373 D18F2x294_dct[1:0]: DRAM Status 4
 374 D18F2x298_dct[1:0]: DRAM Status 5
 374 D18F2x29C_dct[1:0]: DRAM Status 6
 374 D18F3x00: Device/Vendor ID
 374 D18F3x04: Status/Command
 374 D18F3x08: Class Code/Revision ID
 375 D18F3x0C: Header Type

375	D18F3x34: Capability Pointer
375	D18F3x40: MCA NB Control
375	D18F3x44: MCA NB Configuration
379	D18F3x48: MCA NB Status Low
379	D18F3x4C: MCA NB Status High
379	D18F3x50: MCA NB Address Low
379	D18F3x54: MCA NB Address High
379	D18F3x58: Scrub Rate Control
380	D18F3x5C: DRAM Scrub Address Low
381	D18F3x60: DRAM Scrub Address High
381	D18F3x64: Hardware Thermal Control (HTC)
382	D18F3x68: Software P-state Limit
382	D18F3x6C: Data Buffer Count
383	D18F3x70: SRI to XBAR Command Buffer Count
384	D18F3x74: XBAR to SRI Command Buffer Count
385	D18F3x78: MCT to XBAR Buffer Count
386	D18F3x7C: Free List Buffer Count
387	D18F3x[84:80]: ACPI Power State Control
390	D18F3x88: NB Configuration 1 Low (NB_CFG1_LO)
390	D18F3x8C: NB Configuration 1 High (NB_CFG1_HI)
390	D18F3x90: GART Aperture Control
391	D18F3x94: GART Aperture Base
391	D18F3x98: GART Table Base
391	D18F3x9C: GART Cache Control
392	D18F3xA0: Power Control Miscellaneous
393	D18F3xA4: Reported Temperature Control
394	D18F3xA8: Pop Up and Down P-states
394	D18F3xB0: On-Line Spare Control
396	D18F3xB8: NB Array Address
396	D18F3xBC: NB Array Data Port
396	D18F3xBC_x8: DRAM ECC
397	D18F3xC4: SBI P-state Limit
398	D18F3xD4: Clock Power/Timing Control 0
400	D18F3xD8: Clock Power/Timing Control 1
400	D18F3xDC: Clock Power/Timing Control 2
401	D18F3xE4: Thermtrip Status
402	D18F3xE8: Northbridge Capabilities
403	D18F3xFC: CUID Family/Model/Stepping
403	D18F3x138: DCT0 Bad Symbol Identification
404	D18F3x13C: DCT1 Bad Symbol Identification
404	D18F3x140: SRI to XCS Token Count
405	D18F3x144: MCT to XCS Token Count
405	D18F3x1[54,50,4C,48]: Link to XCS Token Count
407	D18F3x158: Link to XCS Token Count
407	D18F3x160: NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)
408	D18F3x168: NB Machine Check Misc (Link Thresholding) 1 (MC4_MISC1)
408	D18F3x170: NB Machine Check Misc (L3 Thresholding) 2 (MC4_MISC2)
409	D18F3x17C: Extended Freelist Buffer Count
409	D18F3x180: Extended NB MCA Configuration
410	D18F3x188: NB Configuration 2 (NB_CFG2)
410	D18F3x190: Downcore Control
411	D18F3x1A0: Core Interface Buffer Count
411	D18F3x1B8: L3 Control 1
413	D18F3x1C4: L3 Cache Parameter
413	D18F3x1CC: IBS Control
413	D18F3x1D4: Probe Filter Control
415	D18F3x1E4: SBI Control
416	D18F3x1E8: SBI Address
416	D18F3x1EC: SBI Data
417	D18F3x1EC_x100: SB-RMI Revision
417	D18F3x1FC: Product Information Register 1
418	D18F4x00: Device/Vendor ID
418	D18F4x04: Status/Command

418 D18F4x08: Class Code/Revision ID
 418 D18F4x0C: Header Type
 419 D18F4x34: Capabilities Pointer
 419 D18F4x[E0,C0,A0,80]: Sublink 1 Capability
 419 D18F4x[E4,C4,A4,84]: Sublink 1 Control
 419 D18F4x[E8,C8,A8,88]: Sublink 1 Frequency/Revision
 419 D18F4x[EC,CC,AC,8C]: Sublink 1 Feature Capability
 420 D18F4x[F0,D0,B0,90]: Sublink 1 Base Channel Buffer Count
 420 D18F4x[F4,D4,B4,94]: Sublink 1 Isochronous Channel Buffer Count
 420 D18F4x[F8,D8,B8,98]: Sublink 1 Link Type
 420 D18F4x[FC,DC,BC,9C]: Sublink 1 Link Frequency Extension
 420 D18F4x104: TDP Accumulator Divisor Control
 421 D18F4x108: TDP Limit 1
 421 D18F4x10C: TDP Limit 2
 421 D18F4x110: Sample and Residency Timers
 421 D18F4x11[C:8]: C-state Control
 424 D18F4x128: C-state Policy Control 1
 425 D18F4x15C: Core Performance Boost Control
 426 D18F4x164: Fixed Errata
 426 D18F4x16C: APM TDP Control
 426 D18F4x1[98,90,88,80]: Link Phy Offset
 428 D18F4x1[9C,94,8C,84]: Link Phy Data Port
 428 D18F4x1[9C,94,8C,84]_x[D0,C0]: Link Phy Impedance
 429 D18F4x1[9C,94,8C,84]_x[D1,C1]: Link Phy Receiver Loop Filter
 431 D18F4x1[9C,94,8C,84]_x[D3,C3]: Link Phy Timing Margin
 432 D18F4x1[9C,94,8C,84]_x[D4,C4]: Link Phy DFE and DFR Control
 433 D18F4x1[9C,94,8C,84]_x[D5,C5]: Link Phy Transmit Control
 434 D18F4x1[9C,94,8C,84]_x[DF,CF]: Link FIFO Read Pointer Optimization
 435 D18F4x1[9C,94,8C,84]_xE0: Link Phy Compensation and Calibration Control 1
 435 D18F4x1[9C,94,8C,84]_xE3: Link Phy PLL Control
 436 D18F4x1[9C,94,8C,84]_x100: Link BIST Control
 437 D18F4x1[9C,94,8C,84]_x104: Link BIST Southbound TX Pattern Control
 438 D18F4x1[9C,94,8C,84]_x108: Link BIST Southbound TX Pattern Buffer 1
 438 D18F4x1[9C,94,8C,84]_x10C: Link BIST Southbound TX Mask
 438 D18F4x1[9C,94,8C,84]_x110: Link BIST Southbound TX Inversion
 438 D18F4x1[9C,94,8C,84]_x114: Link BIST Southbound TX Pattern Buffer 2
 439 D18F4x1[9C,94,8C,84]_x118: Link BIST Southbound TX Pattern Buffer 2 Enable
 439 D18F4x1[9C,94,8C,84]_x11C: Link BIST Southbound TX Pattern Buffer Extension
 439 D18F4x1[9C,94,8C,84]_x120: Link BIST Southbound TX Scramble
 439 D18F4x1[9C,94,8C,84]_x124: Link BIST Northbound RX Pattern Control
 440 D18F4x1[9C,94,8C,84]_x128: Link BIST Northbound RX Pattern Buffer 1
 440 D18F4x1[9C,94,8C,84]_x12C: Link BIST Northbound RX Mask
 441 D18F4x1[9C,94,8C,84]_x130: Link BIST Northbound RX Inversion
 441 D18F4x1[9C,94,8C,84]_x134: Link BIST Northbound RX Pattern Buffer 2
 441 D18F4x1[9C,94,8C,84]_x138: Link BIST Northbound RX Pattern Buffer 2 Enable
 441 D18F4x1[9C,94,8C,84]_x13C: Link BIST Northbound RX Pattern Buffer Extension
 442 D18F4x1[9C,94,8C,84]_x140: Link BIST Northbound RX Scramble
 442 D18F4x1[9C,94,8C,84]_x144: Link BIST Northbound RX Error Status
 442 D18F4x1[9C,94,8C,84]_x148: Link BIST Northbound RX Per-Lane Error Count 1
 443 D18F4x1[9C,94,8C,84]_x14C: Link BIST Northbound RX Per-Lane Error Count 2
 443 D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]6_dm[1]: Link Phy DFE and DFR Control
 444 D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]A_dm[1]: Link Phy DLL Control
 446 D18F4x1[9C,94,8C,84]_x[5:4][9:0][8,0]F_dm[1]: Link Phy Receiver DLL Control and Test 5
 448 D18F4x1[9C,94,8C,84]_x[5:4][9:0][9,1]1_dm[1]: Link Phy Receiver Process Control
 449 D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]0_dm[1]: Link Phy Transmit Control
 450 D18F4x1[9C,94,8C,84]_x6[9:8]84_dm[1]: Link Phy Transmit Clock Phase Control
 451 D18F4x1[9C,94,8C,84]_x[7:6][9:0][8,0]C_dm[1]: Link Phy Tx Deemphasis and Margin Test Control
 454 D18F4x1B8: Processor TDP
 454 D18F4x1C4: L3 Power Control
 455 D18F4x1C8: L3 Hit Statistics
 455 D18F4x1CC: L3 Control 2
 456 D18F4x1D4: Compute Unit Based L3 Cache Partitioning
 457 D18F5x00: Device/Vendor ID

457 D18F5x04: Status/Command
457 D18F5x08: Class Code/Revision ID
457 D18F5x0C: Header Type
457 D18F5x34: Capabilities Pointer
457 D18F5x[70,60,50,40]: Northbridge Performance Event Select Low
458 D18F5x[74,64,54,44]: Northbridge Performance Event Select High
458 D18F5x[78,68,58,48]: Northbridge Performance Event Counter Low
458 D18F5x[7C,6C,5C,4C]: Northbridge Performance Event Counter High
458 D18F5x80: Compute Unit Status
459 D18F5x84: Northbridge Capabilities 2
459 D18F5x88: NB Configuration 4 (NB_CFG4)
460 D18F5xE0: Processor TDP Running Average
460 D18F5xE8: TDP Limit 3
461 D18F5x128: Clock Power/Timing Control 3
461 D18F5x1[6C:60]: Northbridge P-state [3:0]
462 D18F5x170: Northbridge P-state Control
463 D18F5x174: Northbridge P-state Status
464 D18F5x190: Link Product Information
464 D18F5x194: Name String Address Port
464 D18F5x198: Name String Data Port
464 D18F5x198_x[B:0]: Name String Data
465 APIC20: APIC ID
465 APIC30: APIC Version
465 APIC80: Task Priority (TPR)
465 APIC90: Arbitration Priority (APR)
466 APICA0: Processor Priority (PPR)
466 APICB0: End of Interrupt
466 APICC0: Remote Read
466 APICD0: Logical Destination (LDR)
466 APICE0: Destination Format
467 APICF0: Spurious-Interrupt Vector (SVR)
467 APIC[170:100]: In-Service (ISR)
467 APIC[1F0:180]: Trigger Mode (TMR)
468 APIC[270:200]: Interrupt Request (IRR)
468 APIC280: Error Status
469 APIC300: Interrupt Command Low (ICR Low)
470 APIC310: Interrupt Command High (ICR High)
470 APIC320: LVT Timer
471 APIC330: LVT Thermal Sensor
471 APIC340: LVT Performance Monitor
471 APIC3[60:50]: LVT LINT[1:0]
472 APIC370: LVT Error
472 APIC380: Timer Initial Count
472 APIC390: Timer Current Count
472 APIC3E0: Timer Divide Configuration
473 APIC400: Extended APIC Feature
473 APIC410: Extended APIC Control
474 APIC420: Specific End Of Interrupt
474 APIC[4F0:480]: Interrupt Enable
474 APIC[530:500]: Extended Interrupt [3:0] Local Vector Table
475 CPUID Fn0000_0000_EAX: Processor Vendor and Largest Standard Function Number
475 CPUID Fn0000_0000_E[D,C,B]X: Processor Vendor
476 CPUID Fn0000_0001_EAX: Family, Model, Stepping Identifiers
476 CPUID Fn0000_0001_EBX: LocalApicId, LogicalProcessorCount, CLFlush
476 CPUID Fn0000_0001_ECX: Feature Identifiers
477 CPUID Fn0000_0001_EDX: Feature Identifiers
478 CPUID Fn0000_000[4,3,2]: Reserved
478 CPUID Fn0000_0005_EAX: Monitor/MWait
478 CPUID Fn0000_0005_EBX: Monitor/MWait
478 CPUID Fn0000_0005_ECX: Monitor/MWait
478 CPUID Fn0000_0005_EDX: Monitor/MWait
479 CPUID Fn0000_0006_EAX: Thermal and Power Management
479 CPUID Fn0000_0006_EBX: Thermal and Power Management

479 CUID Fn0000_0006_ECX: Thermal and Power Management
 479 CUID Fn0000_0006_EDX: Thermal and Power Management
 479 CUID Fn0000_0007_EAX_x0: Structured Extended Feature Identifiers (ECX=0)
 479 CUID Fn0000_0007_EBX_x0: Structured Extended Feature Identifiers (ECX=0)
 479 CUID Fn0000_0007_ECX_x0: Structured Extended Feature Identifiers (ECX=0)
 480 CUID Fn0000_0007_EDX_x0: Structured Extended Feature Identifiers (ECX=0)
 480 CUID Fn0000_000[C:8]: Reserved
 480 CUID Fn0000_000D_EAX_x0: Processor Extended State Enumeration (ECX=0)
 480 CUID Fn0000_000D_EBX_x0: Processor Extended State Enumeration (ECX=0)
 480 CUID Fn0000_000D_ECX_x0: Processor Extended State Enumeration (ECX=0)
 480 CUID Fn0000_000D_EDX_x0: Processor Extended State Enumeration (ECX=0)
 480 CUID Fn0000_000D_EAX_x2: Processor Extended State Enumeration (ECX=2)
 481 CUID Fn0000_000D_EBX_x2: Processor Extended State Enumeration (ECX=2)
 481 CUID Fn0000_000D_ECX_x2: Processor Extended State Enumeration (ECX=2)
 481 CUID Fn0000_000D_EDX_x2: Processor Extended State Enumeration (ECX=2)
 481 CUID Fn0000_000D_EAX_x3E: Processor Extended State Enumeration (ECX=62)
 481 CUID Fn0000_000D_EBX_x3E: Processor Extended State Enumeration (ECX=62)
 481 CUID Fn0000_000D_ECX_x3E: Processor Extended State Enumeration (ECX=62)
 481 CUID Fn0000_000D_EDX_x3E: Processor Extended State Enumeration (ECX=62)
 482 CUID Fn8000_0000_EAX: Largest Extended Function Number
 482 CUID Fn8000_0000_E[D,C,B]X: Processor Vendor
 482 CUID Fn8000_0001_EAX: Family, Model, Stepping Identifiers
 482 CUID Fn8000_0001_EBX: BrandId Identifier
 483 CUID Fn8000_0001_ECX: Feature Identifiers
 484 CUID Fn8000_0001_EDX: Feature Identifiers
 485 CUID Fn8000_000[4:2]_E[D,C,B,A]X: Processor Name String Identifier
 485 CUID Fn8000_0005_EAX: L1 TLB 2M/4M Identifiers
 485 CUID Fn8000_0005_EBX: L1 TLB 4K Identifiers
 486 CUID Fn8000_0005_ECX: L1 Data Cache Identifiers
 486 CUID Fn8000_0005_EDX: L1 Instruction Cache Identifiers
 486 CUID Fn8000_0006_EAX: L2 TLB 2M/4M Identifiers
 487 CUID Fn8000_0006_EBX: L2 TLB 4K Identifiers
 487 CUID Fn8000_0006_ECX: L2 Cache Identifiers
 488 CUID Fn8000_0006_EDX: L3 Cache Identifiers
 488 CUID Fn8000_0007_E[C,B,A]X: Advanced Power Management Information
 488 CUID Fn8000_0007_EDX: Advanced Power Management Information
 489 CUID Fn8000_0008_EAX: Long Mode Address Size Identifiers
 489 CUID Fn8000_0008_EBX: Reserved
 489 CUID Fn8000_0008_ECX: APIC ID Size and Core Count
 490 CUID Fn8000_0008_EDX: Reserved
 490 CUID Fn8000_0009: Reserved
 490 CUID Fn8000_000A_EAX: SVM Revision
 490 CUID Fn8000_000A_EBX: SVM Revision and Feature Identification
 490 CUID Fn8000_000A_ECX: Reserved
 490 CUID Fn8000_000A_EDX: SVM Feature Identification
 491 CUID Fn8000_00[18:0B]: Reserved
 491 CUID Fn8000_0019_EAX: L1 TLB 1G Identifiers
 491 CUID Fn8000_0019_EBX: L2 TLB 1G Identifiers
 492 CUID Fn8000_0019_E[D,C]X: Reserved
 492 CUID Fn8000_001A_EAX: Performance Optimization Identifiers
 492 CUID Fn8000_001A_E[D,C,B]X: Reserved
 492 CUID Fn8000_001B_EAX: Instruction Based Sampling Identifiers
 493 CUID Fn8000_001B_E[D,C,B]X: Instruction Based Sampling Identifiers
 493 CUID Fn8000_001C_EAX: Lightweight Profiling Capabilities 0
 493 CUID Fn8000_001C_EBX: Lightweight Profiling Capabilities 0
 494 CUID Fn8000_001C_ECX: Lightweight Profiling Capabilities 0
 494 CUID Fn8000_001C_EDX: Lightweight Profiling Capabilities 0
 494 CUID Fn8000_001D_EAX_x0: Cache Properties
 495 CUID Fn8000_001D_EAX_x1: Cache Properties
 496 CUID Fn8000_001D_EAX_x2: Cache Properties
 496 CUID Fn8000_001D_EAX_x3: Cache Properties
 497 CUID Fn8000_001D_EAX_x4: Cache Properties
 497 CUID Fn8000_001D_EBX_x0: Cache Properties

497 CUID Fn8000_001D_EBX_x1: Cache Properties
 497 CUID Fn8000_001D_EBX_x2: Cache Properties
 498 CUID Fn8000_001D_EBX_x3: Cache Properties
 498 CUID Fn8000_001D_EBX_x4: Cache Properties
 498 CUID Fn8000_001D_ECX_x0: Cache Properties
 499 CUID Fn8000_001D_ECX_x1: Cache Properties
 499 CUID Fn8000_001D_ECX_x2: Cache Properties
 499 CUID Fn8000_001D_ECX_x3: Cache Properties
 499 CUID Fn8000_001D_ECX_x4: Cache Properties
 499 CUID Fn8000_001D_EDX_x0: Cache Properties
 500 CUID Fn8000_001D_EDX_x1: Cache Properties
 500 CUID Fn8000_001D_EDX_x2: Cache Properties
 500 CUID Fn8000_001D_EDX_x3: Cache Properties
 500 CUID Fn8000_001D_EDX_x4: Cache Properties
 500 CUID Fn8000_001E_EAX: Extended APIC ID
 501 CUID Fn8000_001E_EBX: Compute Unit Identifiers
 501 CUID Fn8000_001E_ECX: Node Identifiers
 501 CUID Fn8000_001E_EDX: Reserved
 501 MSR0000_0000: Load-Store MCA Address
 501 MSR0000_0001: Load-Store MCA Status
 502 MSR0000_0010: Time Stamp Counter (TSC)
 502 MSR0000_001B: APIC Base Address (APIC_BAR)
 502 MSR0000_002A: Cluster ID (EBL_CR_POWERON)
 503 MSR0000_00E7: Max Performance Frequency Clock Count (MPERF)
 503 MSR0000_00E8: Actual Performance Frequency Clock Count (APERF)
 503 MSR0000_00FE: MTRR Capabilities (MTRRcap)
 503 MSR0000_0174: SYSENTER CS (SYSENTER_CS)
 503 MSR0000_0175: SYSENTER ESP (SYSENTER_ESP)
 504 MSR0000_0176: SYSENTER EIP (SYSENTER_EIP)
 504 MSR0000_0179: Global Machine Check Capabilities (MCG_CAP)
 504 MSR0000_017A: Global Machine Check Status (MCG_STAT)
 504 MSR0000_017B: Global Machine Check Exception Reporting Control (MCG_CTL)
 505 MSR0000_01D9: Debug Control (DBG_CTL_MSR)
 505 MSR0000_01DB: Last Branch From IP (BR_FROM)
 505 MSR0000_01DC: Last Branch To IP (BR_TO)
 505 MSR0000_01DD: Last Exception From IP
 506 MSR0000_01DE: Last Exception To IP
 506 MSR0000_020[F:0]: Variable-Size MTRRs Base/Mask
 507 MSR0000_02[6F:68,59:58,50]: Fixed-Size MTRRs
 509 MSR0000_0277: Page Attribute Table (PAT)
 510 MSR0000_02FF: MTRR Default Memory Type (MTRRdefType)
 510 MSR0000_0400: LS Machine Check Control (MC0_CTL)
 511 MSR0000_0401: LS Machine Check Status (MC0_STATUS)
 513 MSR0000_0402: LS Machine Check Address (MC0_ADDR)
 514 MSR0000_0403: LS Machine Check Miscellaneous (MC0_MISC)
 515 MSR0000_0404: IF Machine Check Control (MC1_CTL)
 515 MSR0000_0405: IF Machine Check Status (MC1_STATUS)
 520 MSR0000_0406: IF Machine Check Address (MC1_ADDR)
 522 MSR0000_0407: IF Machine Check Miscellaneous (MC1_MISC)
 522 MSR0000_0408: CU Machine Check Control (MC2_CTL)
 522 MSR0000_0409: CU Machine Check Status (MC2_STATUS)
 527 MSR0000_040A: CU Machine Check Address (MC2_ADDR)
 529 MSR0000_040B: CU Machine Check Miscellaneous (MC2_MISC)
 529 MSR0000_040C: MC3 Machine Check Control (MC3_CTL)
 529 MSR0000_040D: MC3 Machine Check Status (MC3_STATUS)
 530 MSR0000_040E: MC3 Machine Check Address (MC3_ADDR)
 530 MSR0000_040F: MC3 Machine Check Miscellaneous (MC3_MISC)
 530 MSR0000_0410: NB Machine Check Control (MC4_CTL)
 532 MSR0000_0411: NB Machine Check Status (MC4_STATUS)
 537 MSR0000_0412: NB Machine Check Address (MC4_ADDR)
 542 MSR0000_0413: NB Machine Check Misc (DRAM Thresholding) 0 (MC4_MISC0)
 543 MSR0000_0414: EX Machine Check Control (MC5_CTL)
 544 MSR0000_0415: EX Machine Check Status (MC5_STATUS)

546 MSR0000_0416: EX Machine Check Address (MC5_ADDR)
 547 MSR0000_0417: EX Machine Check Miscellaneous (MC5_MISC)
 547 MSR0000_0418: FP Machine Check Control (MC6_CTL)
 547 MSR0000_0419: FP Machine Check Status (MC6_STATUS)
 548 MSR0000_041A: FP Machine Check Address (MC6_ADDR)
 549 MSR0000_041B: FP Machine Check Miscellaneous (MC6_MISC)
 549 MSRC000_0080: Extended Feature Enable (EFER)
 550 MSRC000_0081: SYSCALL Target Address (STAR)
 550 MSRC000_0082: Long Mode SYSCALL Target Address (STAR64)
 550 MSRC000_0083: Compatibility Mode SYSCALL Target Address (STARCOMPAT)
 550 MSRC000_0084: SYSCALL Flag Mask (SYSCALL_FLAG_MASK)
 550 MSRC000_00E7: Read-Only Max Performance Frequency Clock Count (MPerfReadOnly)
 551 MSRC000_00E8: Read-Only Actual Performance Frequency Clock Count (APerfReadOnly)
 551 MSRC000_0100: FS Base (FS_BASE)
 551 MSRC000_0101: GS Base (GS_BASE)
 551 MSRC000_0102: Kernel GS Base (KernelGSbase)
 552 MSRC000_0103: Auxiliary Time Stamp Counter (TSC_AUX)
 552 MSRC000_0104: Time Stamp Counter Ratio (TscRateMsr)
 553 MSRC000_0105: Lightweight Profile Configuration (LWP_CFG)
 553 MSRC000_0106: Lightweight Profile Control Block Address (LWP_CBADDR)
 553 MSRC000_0408: Machine Check Misc 4 (Link Thresholding) 1 (MC4_MISC1)
 554 MSRC000_0409: Machine Check Misc 4 (L3 Thresholding) 2 (MC4_MISC2)
 554 MSRC000_040[F:A]: Reserved
 555 MSRC001_00[03:00]: Performance Event Select (PERF_CTL[3:0])
 555 MSRC001_00[07:04]: Performance Event Counter (PERF_CTR[3:0])
 555 MSRC001_0010: System Configuration (SYS_CFG)
 556 MSRC001_0015: Hardware Configuration (HWCR)
 558 MSRC001_00[18,16]: IO Range Base (IORR_BASE[1:0])
 559 MSRC001_00[19,17]: IO Range Mask (IORR_MASK[1:0])
 559 MSRC001_001A: Top Of Memory (TOP_MEM)
 559 MSRC001_001D: Top Of Memory 2 (TOM2)
 560 MSRC001_001F: NB Configuration 1 (NB_CFG1)
 560 MSRC001_0022: Machine Check Exception Redirection
 561 MSRC001_00[35:30]: Processor Name String
 561 MSRC001_003E: Hardware Thermal Control (HTC)
 562 MSRC001_0044: LS Machine Check Control Mask (MC0_CTL_MASK)
 562 MSRC001_0045: IF Machine Check Control Mask (MC1_CTL_MASK)
 563 MSRC001_0046: CU Machine Check Control Mask (MC2_CTL_MASK)
 563 MSRC001_0047: Reserved (MC3_CTL_MASK)
 563 MSRC001_0048: NB Machine Check Control Mask (MC4_CTL_MASK)
 564 MSRC001_0049: EX Machine Check Control Mask (MC5_CTL_MASK)
 564 MSRC001_004A: FP Machine Check Control Mask (MC6_CTL_MASK)
 565 MSRC001_00[53:50]: IO Trap (SMI_ON_IO_TRAP_[3:0])
 565 MSRC001_0054: IO Trap Control (SMI_ON_IO_TRAP_CTL_STS)
 566 MSRC001_0055: Interrupt Pending
 567 MSRC001_0056: SMI Trigger IO Cycle
 567 MSRC001_0058: MMIO Configuration Base Address
 568 MSRC001_0060: BIST Results
 568 MSRC001_0061: P-state Current Limit
 569 MSRC001_0062: P-state Control
 569 MSRC001_0063: P-state Status
 569 MSRC001_00[6B:64]: P-state [7:0]
 571 MSRC001_0070: COFVID Control
 572 MSRC001_0071: COFVID Status
 573 MSRC001_0072: SBI P-state Limit
 573 MSRC001_0073: C-state Base Address
 574 MSRC001_0075: APLM TDP Limit
 574 MSRC001_0077: Processor Power in TDP
 575 MSRC001_0078: Power Averaging Period
 575 MSRC001_0079: DRAM Controller Command Throttle
 575 MSRC001_0111: SMM Base Address (SMM_BASE)
 576 MSRC001_0112: SMM TSeg Base Address (SMMAddr)
 576 MSRC001_0113: SMM TSeg Mask (SMMMask)

577 MSRC001_0114: Virtual Machine Control (VM_CR)
578 MSRC001_0115: IGNNE
578 MSRC001_0116: SMM Control (SMM_CTL)
579 MSRC001_0117: Virtual Machine Host Save Physical Address (VM_HSAVE_PA)
579 MSRC001_0118: SVM Lock Key
579 MSRC001_011A: Local SMI Status
579 MSRC001_0140: OS Visible Work-around MSR0 (OSVW_ID_Length)
579 MSRC001_0141: OS Visible Work-around MSR1 (OSVW_Status)
580 MSRC001_020[A,8,6,4,2,0]: Performance Event Select (PERF_CTL[5:0])
581 MSRC001_020[B,9,7,5,3,1]: Performance Event Counter (PERF_CTR[5:0])
581 MSRC001_024[6,4,2,0]: Northbridge Performance Event Select (NB_PERF_CTL[3:0])
583 MSRC001_024[7,5,3,1]: Northbridge Performance Event Counter (NB_PERF_CTR[3:0])
584 MSRC001_1003: Thermal and Power Management CPUID Features
584 MSRC001_1004: CPUID Features (Features)
585 MSRC001_1005: Extended CPUID Features (ExtFeatures)
587 MSRC001_1020: Load-Store Configuration (LS_CFG)
587 MSRC001_1021: Instruction Cache Configuration (IC_CFG)
588 MSRC001_1022: Data Cache Configuration (DC_CFG)
588 MSRC001_1023: Combined Unit Configuration (CU_CFG)
589 MSRC001_1028: Floating Point Configuration (FP_CFG)
590 MSRC001_1029: Decode Configuration (DE_CFG)
590 MSRC001_102A: Combined Unit Configuration 2 (CU_CFG2)
591 MSRC001_102B: Combined Unit Configuration 3 (CU_CFG3)
592 MSRC001_102C: Execution Unit Configuration (EX_CFG)
592 MSRC001_102D: Load-Store Configuration 2 (LS_CFG2)
593 MSRC001_1030: IBS Fetch Control (IC_IBS_CTL)
594 MSRC001_1031: IBS Fetch Linear Address (IC_IBS_LIN_AD)
594 MSRC001_1032: IBS Fetch Physical Address (IC_IBS_PHYS_AD)
595 MSRC001_1033: IBS Execution Control (SC_IBS_CTL)
596 MSRC001_1034: IBS Op Logical Address (IBSOP_RIP)
596 MSRC001_1035: IBS Op Data (SC_IBS_DATA)
596 MSRC001_1036: IBS Op Data 2 (IbsOpData2)
597 MSRC001_1037: IBS Op Data 3 (DC_IBS_DATA, IbsOpData3)
598 MSRC001_1038: IBS DC Linear Address (DC_IBS_LIN_ADDR)
598 MSRC001_1039: IBS DC Physical Address (DC_IBS_PHYS_ADDR)
598 MSRC001_103A: IBS Control
599 MSRC001_103B: IBS Branch Target Address (BP_IBSTGT_RIP)
599 PMCx000: FPU Pipe Assignment
599 PMCx001: FP Scheduler Empty
600 PMCx003: Retired Floating Point Ops
600 PMCx004: Number of Move Elimination and Scalar Op Optimization
600 PMCx005: Retired Serializing Ops
600 PMCx006: Number of Cycles that a Bottom-Execute uop is in the FP Scheduler
601 PMCx020: Segment Register Loads
601 PMCx021: Pipeline Restart Due to Self-Modifying Code
601 PMCx022: Pipeline Restart Due to Probe Hit
601 PMCx023: Load Queue/Store Queue Full
601 PMCx024: Locked Operations
602 PMCx026: Retired CLFLUSH Instructions
602 PMCx027: Retired CPUID Instructions
602 PMCx029: LS Dispatch
602 PMCx02A: Canceled Store to Load Forward Operations
602 PMCx02B: SMIs Received
602 PMCx030: Executed CLFLUSH Instructions
602 PMCx032: Misaligned Stores
603 PMCx034: FP +Load Buffer Stall
603 PMCx040: Data Cache Accesses
603 PMCx041: Data Cache Misses
603 PMCx042: Data Cache Refills from L2 or System
603 PMCx043: Data Cache Refills from System
603 PMCx045: Unified TLB Hit
604 PMCx046: Unified TLB Miss
604 PMCx047: Misaligned Accesses

604 PMCx04B: Prefetch Instructions Dispatched
605 PMCx052: Ineffective Software Prefetches
605 PMCx065: Memory Requests by Type
605 PMCx067: Data Prefetcher
605 PMCx068: MAB Requests
606 PMCx069: MAB Wait Cycles
606 PMCx06C: Response From System on Cache Refills
606 PMCx06D: Octwords Written to System
607 PMCx075: Cache Cross-invalidates
607 PMCx076: CPU Clocks not Halted
607 PMCx07D: Requests to L2 Cache
607 PMCx07E: L2 Cache Misses
608 PMCx07F: L2 Fill/Writeback
608 PMCx165: Page Splintering
608 PMCx16C: L2 Prefetcher Trigger Events
608 PMCx080: Instruction Cache Fetches
609 PMCx081: Instruction Cache Misses
609 PMCx082: Instruction Cache Refills from L2
609 PMCx083: Instruction Cache Refills from System
609 PMCx084: L1 ITLB Miss, L2 ITLB Hit
609 PMCx085: L1 ITLB Miss, L2 ITLB Miss
609 PMCx086: Pipeline Restart Due to Instruction Stream Probe
609 PMCx087: Instruction Fetch Stall
609 PMCx088: Return Stack Hits
610 PMCx089: Return Stack Overflows
610 PMCx08B: Instruction Cache Victims
610 PMCx08C: Instruction Cache Lines Invalidated
610 PMCx099: ITLB Reloads
610 PMCx09A: ITLB Reloads Aborted
610 PMCx0C0: Retired Instructions
611 PMCx0C1: Retired uops
611 PMCx0C2: Retired Branch Instructions
611 PMCx0C3: Retired Mispredicted Branch Instructions
611 PMCx0C4: Retired Taken Branch Instructions
611 PMCx0C5: Retired Taken Branch Instructions Mispredicted
611 PMCx0C6: Retired Far Control Transfers
611 PMCx0C7: Retired Branch Resyncs
611 PMCx0C8: Retired Near Returns
611 PMCx0C9: Retired Near Returns Mispredicted
611 PMCx0CA: Retired Indirect Branches Mispredicted
612 PMCx0CB: Retired MMX/FP Instructions
612 PMCx0CD: Interrupts-Masked Cycles
612 PMCx0CE: Interrupts-Masked Cycles with Interrupt Pending
612 PMCx0CF: Interrupts Taken
612 PMCx0D0: Decoder Empty
612 PMCx0D1: Dispatch Stalls
613 PMCx0D3: Microsequencer Stall due to Serialization
613 PMCx0D5: Dispatch Stall for Instruction Retire Q Full
613 PMCx0D6: Dispatch Stall for Integer Scheduler Queue Full
613 PMCx0D7: Dispatch Stall for FP Scheduler Queue Full
613 PMCx0D8: Dispatch Stall for LDQ Full
613 PMCx0D9: Microsequencer Stall Waiting for All Quiet
613 PMCx0DB: FPU Exceptions
614 PMCx0D[F:C]: DR[3:0] Breakpoint Matches
614 PMCx1C0: Retired x87 Floating Point Operations
614 PMCx1CF: Tagged IBS Ops
615 PMCx1D8: Dispatch Stall for STQ Full
615 PMCx1DD: Cycles Without Dispatch Due To Integer PRF Tokens
615 PMCx1DE: Cycles Without Dispatch Due to FP PRF Tokens
615 NBPMCx0E0: DRAM Accesses
616 NBPMCx0E1: DRAM Controller Page Table Overflows
616 NBPMCx0E2: Memory Controller DRAM Command Slots Missed
616 NBPMCx0E3: Memory Controller Turnarounds

617 NBPMCx0E4: Memory Controller Bypass Counter Saturation
617 NBPMCx0E8: Thermal Status
617 NBPMCx0E9: CPU/IO Requests to Memory/IO
618 NBPMCx0EA: Cache Block Commands
619 NBPMCx0EB: Sized Commands
619 NBPMCx0EC: Probe Responses and Upstream Requests
620 NBPMCx0EE: GART Events
620 NBPMCx[1F9,0F8,0F7,0F6]: Link Transmit Bandwidth
621 NBPMCx1E0: CPU to DRAM Requests to Target Node
621 NBPMCx1E1: IO to DRAM Requests to Target Node
621 NBPMCx1E2: CPU Read Command Latency to Target Node 0-3
622 NBPMCx1E3: CPU Read Command Requests to Target Node 0-3
622 NBPMCx1E4: CPU Read Command Latency to Target Node 4-7
623 NBPMCx1E5: CPU Read Command Requests to Target Node 4-7
623 NBPMCx1E6: CPU Command Latency to Target Node 0-3/4-7
623 NBPMCx1E7: CPU Requests to Target Node 0-3/4-7
624 NBPMCx1EA: Request Cache Status 0
624 NBPMCx1EB: Request Cache Status 1
625 NBPMCx1F0: Memory Controller Requests
625 NBPMCx4E0: Read Request to L3 Cache
626 NBPMCx4E1: L3 Cache Misses
626 NBPMCx4E2: L3 Fills caused by L2 Evictions
626 NBPMCx4E3: L3 Evictions
627 NBPMCx4ED: Non-canceled L3 Read Requests
627 NBPMCx4EF: L3 Latency