

e200z0 Power Architecture™ Core Reference Manual

Supports

e200z0

e200z0h

e200z0CORERM

Rev. 0

4/2008

How to Reach Us:

Home Page:

www.freescale.com

Web Support:

<http://www.freescale.com/support>

USA/Europe or Locations Not Listed:

Freescale Semiconductor, Inc.
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or
+1-480-768-2130
www.freescale.com/support

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064
Japan
0120 191014 or
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
+1-800 441-2447 or
+1-303-675-2140
Fax: +1-303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org. IEEE ISTO 5001 and 1149.1 are registered trademarks of the Institute of Electrical and Electronics Engineers, Inc. (IEEE). This product is not endorsed or approved by the IEEE. All other product or service names are the property of their respective owners.

© Freescale Semiconductor, Inc., 2008. All rights reserved.





e200z0 and e200z0h Overview	1
Register Model	2
Instruction Model	3
Instruction Pipeline and Execution Timing	4
Interrupts and Exceptions	5
Core Complex Interfaces	6
Power Management	7
Debug Support	8
Nexus 2+ Module	9
Register Summary	A
Glossary	GLO
Index	IND



1	e200z0 and e200z0h Overview
2	Register Model
3	Instruction Model
4	Instruction Pipeline and Execution Timing
5	Interrupts and Exceptions
6	Core Complex Interfaces
7	Power Management
8	Debug Support
9	Nexus 2+ Module
A	Register Summary
GLO	Glossary
IND	Index

Contents

Paragraph Number	Title	Page Number
------------------	-------	-------------

Chapter 1 e200z0 and e200z0h Overview

1.1	Overview of the e200z0 and e200z0h Cores	1-1
1.1.1	Features	1-1
1.1.2	Microarchitecture Summary	1-2
1.1.2.1	Instruction Unit Features	1-5
1.1.2.2	Integer Unit Features	1-6
1.1.2.3	Load/Store Unit Features	1-6
1.1.2.4	e200z0h System Bus Features	1-6
1.1.2.5	e200z0 System Bus Features	1-6
1.1.2.6	Nexus 2+ Features	1-7

Chapter 2 Register Model

2.1	Power Architecture Book E Registers	2-4
2.2	e200-Specific Special Purpose Registers.....	2-6
2.3	Special Purpose Register Descriptions	2-7
2.3.1	Machine State Register (MSR)	2-7
2.3.2	Processor ID Register (PIR)	2-9
2.3.3	Processor Version Register (PVR).....	2-10
2.3.4	System Version Register (SVR).....	2-10
2.3.5	Integer Exception Register (XER).....	2-11
2.3.6	Exception Syndrome Register	2-12
2.3.6.1	Power Architecture VLE Mode Instruction Syndrome	2-13
2.3.6.2	Misaligned Instruction Fetch Syndrome.....	2-13
2.3.6.3	Precise External Termination Error Syndrome.....	2-14
2.3.7	Machine Check Syndrome Register (MCSR).....	2-14
2.3.8	Debug Registers.....	2-16
2.3.9	Hardware Implementation Dependent Register 0 (HID0).....	2-16
2.3.10	Hardware Implementation Dependent Register 1 (HID1).....	2-18
2.3.11	Branch Unit Control and Status Register (BUCSR).....	2-19
2.3.12	L1 Cache Configuration Register (L1CFG0)	2-20
2.3.13	MMU Configuration Register (MMUCFG)	2-20
2.4	SPR Register Access.....	2-20
2.4.1	Invalid SPR References	2-20
2.4.2	Synchronization Requirements for SPRs.....	2-21
2.4.3	Special Purpose Register Summary	2-21
2.4.4	Reset Settings.....	2-23

Contents

Paragraph Number	Title	Page Number
------------------	-------	-------------

Chapter 3 Instruction Model

3.1	Unsupported Instructions and Instruction Forms.....	3-1
3.2	Optionally Supported Instructions and Instruction Forms.....	3-1
3.3	Memory Access Alignment Support.....	3-2
3.4	Memory Synchronization and Reservation Instructions.....	3-2
3.5	Branch Prediction	3-3
3.6	Interruption of Instructions by Interrupt Requests.....	3-3
3.7	New e200 Instructions	3-3
3.7.1	ISEL APU	3-3
3.7.2	Debug APU.....	3-4
3.7.3	WAIT APU	3-5
3.8	Unimplemented SPRs and Read-Only SPRs	3-6
3.9	Invalid Forms of Instructions.....	3-6
3.9.1	Load and Store with Update Instructions	3-6
3.9.2	Load Multiple Word (e_lmw) Instruction.....	3-6
3.9.3	Instructions with Reserved Fields Non-Zero	3-7
3.10	Optionally Supported APU Instructions	3-7

Chapter 4 Instruction Pipeline and Execution Timing

4.1	Overview of Operation	4-1
4.1.1	Control Unit	4-4
4.1.2	Instruction Unit.....	4-4
4.1.3	Branch Unit.....	4-4
4.1.4	Instruction Decode Unit.....	4-4
4.1.5	Exception Handling	4-4
4.2	Execution Units.....	4-4
4.2.1	Integer Execution Unit.....	4-4
4.2.2	Load/Store Unit.....	4-5
4.3	Instruction Pipeline	4-5
4.3.1	Description of Pipeline Stages	4-6
4.3.2	Instruction Buffers	4-7
4.3.2.1	Branch Prediction in e200z0h.....	4-7
4.3.3	Single-Cycle Instruction Pipeline Operation	4-9
4.3.4	Basic Load and Store Instruction Pipeline Operation.....	4-9
4.3.5	Change-of-Flow Instruction Pipeline Operation.....	4-11
4.3.6	Basic Multi-Cycle Instruction Pipeline Operation.....	4-12
4.3.7	Additional Examples of Instruction Pipeline Operation for Load and Store.....	4-12

Contents

Paragraph Number	Title	Page Number
4.3.8	Move to/from SPR Instruction Pipeline Operation.....	4-13
4.4	Control Hazards	4-14
4.5	Instruction Serialization	4-14
4.5.1	Completion Serialization	4-14
4.5.2	Dispatch Serialization	4-15
4.5.3	Refetch Serialization.....	4-15
4.6	Interrupt Recognition and Exception Processing.....	4-15
4.7	Instruction Timings	4-19
4.8	Operand Placement on Performance.....	4-20

Chapter 5 Interrupts and Exceptions

5.1	e200 Interrupts	5-2
5.2	Exception Syndrome Register	5-3
5.3	Machine State Register	5-5
5.4	Machine Check Syndrome Register (MCSR).....	5-7
5.5	Interrupt Vector Prefix Register (IVPR)	5-8
5.6	Interrupt Vector Offset Values (IVORxx)	5-9
5.7	Interrupt Definitions	5-10
5.7.1	Critical Input Interrupt (IVOR0).....	5-10
5.7.2	Machine Check Interrupt (IVOR1).....	5-11
5.7.2.1	Machine Check Interrupt Enabled (MSR[ME]=1)	5-11
5.7.2.2	Checkstop State	5-12
5.7.2.3	Non-Maskable Interrupts (NMI) in e200z0h.....	5-13
5.7.3	Data Storage Interrupt (IVOR2)	5-13
5.7.4	Instruction Storage Interrupt (IVOR3)	5-14
5.7.5	External Input Interrupt (IVOR4).....	5-14
5.7.6	Alignment Interrupt (IVOR5).....	5-15
5.7.7	Program Interrupt (IVOR6)	5-15
5.7.8	System Call Interrupt (IVOR8).....	5-17
5.7.9	Auxiliary Processor Unavailable Interrupt (IVOR9).....	5-17
5.7.10	Debug Interrupt (IVOR15)	5-17
5.7.11	System Reset Interrupt.....	5-20
5.8	Exception Recognition and Priorities	5-21
5.8.1	Exception Priorities.....	5-23
5.9	Interrupt Processing	5-25
5.9.1	Enabling and Disabling Exceptions.....	5-26
5.9.2	Returning from an Interrupt Handler	5-27
5.10	Process Switching.....	5-28

Contents

Paragraph Number	Title	Page Number
Chapter 6		
Core Complex Interfaces		
6.1	Signal Index	6-2
6.2	Internal Interface Signals	6-8
6.3	Signal Descriptions	6-10
6.3.1	e200 Processor Clock (<i>m_clk</i>)	6-10
6.3.2	Reset-Related Signals	6-10
6.3.2.1	Power-On Reset (<i>m_por</i>)	6-11
6.3.2.2	Reset (<i>p_reset_b</i>)	6-11
6.3.2.3	Reset Out (<i>p_resetout_b</i>)	6-11
6.3.2.4	Reset Base (<i>p_rstbase</i> [0:29])	6-11
6.3.2.5	JTAG/OnCE Reset (<i>j_trst_b</i>)	6-11
6.3.3	Address and Data Buses	6-11
6.3.3.1	Address Bus (<i>p_d_haddr</i> [31:0], <i>p_i_haddr</i> [31:0])	6-12
6.3.3.2	Read Data Bus (<i>p_d_hrdata</i> [31:0], <i>p_i_hrdata</i> [31:0])	6-12
6.3.3.3	Write Data Bus (<i>p_d_hwdata</i> [31:0])	6-12
6.3.4	Transfer Attribute Signals	6-12
6.3.4.1	Transfer Type (<i>p_d_htrans</i> [1:0], <i>p_i_htrans</i> [1:0])	6-13
6.3.4.2	Write (<i>p_d_hwrite</i> , <i>p_i_hwrite</i>)	6-13
6.3.4.3	Transfer Size (<i>p_d_hsize</i> [1:0], <i>p_i_hsize</i> [1:0])	6-13
6.3.4.4	Burst Type (<i>p_d_hburst</i> [2:0], <i>p_i_hburst</i> [2:0])	6-14
6.3.4.5	Protection Control (<i>p_d_hprot</i> [5:0], <i>p_i_hprot</i> [5:0])	6-14
6.3.5	Byte Lane Specification	6-15
6.3.5.1	Unaligned Access (<i>p_d_hunalign</i> , <i>p_i_hunalign</i>)	6-15
6.3.5.2	Byte Strobes (<i>p_d_hbstrb</i> [3:0], <i>p_i_hbstrb</i> [3:0])	6-15
6.3.6	Transfer Control Signals	6-18
6.3.6.1	Transfer Ready (<i>p_d_hready</i> , <i>p_i_hready</i>)	6-18
6.3.6.2	Transfer Response (<i>p_d_hresp</i> [2:0], <i>p_i_hresp</i> [1:0])	6-18
6.3.7	Interrupt Signals	6-19
6.3.7.1	External Input Interrupt Request (<i>p_extint_b</i>)	6-19
6.3.7.2	Critical Input Interrupt Request (<i>p_critint_b</i>)	6-19
6.3.7.3	Non-Maskable Input Interrupt Request (<i>p_nmi_b</i>) on e200z0h	6-19
6.3.7.4	Interrupt Pending (<i>p_ipend</i>)	6-20
6.3.7.5	Autovector (<i>p_avec_b</i>)	6-20
6.3.7.6	Interrupt Vector Offset (<i>p_voffset</i> [0:9])	6-20
6.3.7.7	Interrupt Vector Acknowledge (<i>p_iack</i>)	6-20
6.3.7.8	Machine Check (<i>p_mcp_b</i>)	6-21
6.3.8	Processor Reservation Signals	6-21
6.3.8.1	CPU Reservation Status (<i>p_rsrv</i>)	6-21
6.3.8.2	CPU Reservation Clear (<i>p_rsrv_clr</i>)	6-21

Contents

Paragraph Number	Title	Page Number
6.3.9	Miscellaneous Processor Signals	6-22
6.3.9.1	PID0 Outputs (<i>p_pid0[0:7]</i>).....	6-22
6.3.9.2	PID0 Update (<i>p_pid0_updt</i>)	6-22
6.3.9.3	HID1 System Control (<i>p_hid1_sysctl[0:7]</i>).....	6-22
6.3.10	Processor State Signals	6-22
6.3.10.1	Processor Status (<i>p_pstat[0:5]</i>).....	6-22
6.3.10.2	Processor Exception Enable MSR Values (<i>p_EE, p_CE, p_DE, p_ME</i>).....	6-23
6.3.10.3	Processor Machine Check (<i>p_mcp_out</i>)	6-24
6.3.10.4	Processor Checkstop (<i>p_chkstop</i>).....	6-24
6.3.11	Power Management Control Signals	6-24
6.3.11.1	Processor Waiting (<i>p_waiting</i>).....	6-24
6.3.11.2	Processor Halt Request (<i>p_halt</i>).....	6-24
6.3.11.3	Processor Halted (<i>p_halted</i>)	6-24
6.3.11.4	Processor Stop Request (<i>p_stop</i>)	6-24
6.3.11.5	Processor Stopped (<i>p_stopped</i>).....	6-24
6.3.11.6	Low-Power Mode Signals (<i>p_doze, p_nap, p_sleep</i>)	6-24
6.3.11.7	Wakeup (<i>p_wakeup</i>)	6-25
6.3.12	Debug Event Signals.....	6-25
6.3.12.1	Unconditional Debug Event (<i>p_ude</i>)	6-25
6.3.12.2	External Debug Event 1 (<i>p_devt1</i>)	6-25
6.3.12.3	External Debug Event 2 (<i>p_devt2</i>)	6-26
6.3.13	Debug/Emulation (Nexus 1/OnCE) Support Signals.....	6-26
6.3.13.1	OnCE Enable (<i>jd_en_once</i>)	6-26
6.3.13.2	Debug Session (<i>jd_debug_b</i>).....	6-27
6.3.13.3	Debug Request (<i>jd_de_b</i>)	6-27
6.3.13.4	DE_b Active High Output Enable (<i>jd_de_en</i>).....	6-27
6.3.13.5	Processor Clock On (<i>jd_mclk_on</i>)	6-27
6.3.13.6	Watchpoint Events (<i>jd_watchpoint[0:5]</i>).....	6-27
6.3.14	Development Support (Nexus2+) Signals	6-28
6.3.15	JTAG Support Signals	6-28
6.3.15.1	JTAG/OnCE Serial Input (<i>j_tdi</i>)	6-28
6.3.15.2	JTAG/OnCE Serial Clock (<i>j_tclk</i>)	6-28
6.3.15.3	JTAG/OnCE Serial Output (<i>j_tdo</i>)	6-29
6.3.15.4	JTAG/OnCE Test Mode Select (<i>j_tms</i>).....	6-29
6.3.15.5	JTAG/OnCE Test Reset (<i>j_trst_b</i>)	6-29
6.3.15.6	Test-Logic-Reset (<i>j_tst_log_rst</i>).....	6-30
6.3.15.7	Run-Test/Idle (<i>j_rti</i>).....	6-30
6.3.15.8	Capture IR (<i>j_capture_ir</i>)	6-30
6.3.15.9	Shift IR (<i>j_shift_ir</i>)	6-30
6.3.15.10	Update IR (<i>j_update_ir</i>)	6-30
6.3.15.11	Capture DR (<i>j_capture_dr</i>).....	6-30

Contents

Paragraph Number	Title	Page Number
6.3.15.12	Shift DR (<i>j_shift_dr</i>).....	6-30
6.3.15.13	Update DR (<i>j_update_gp_reg</i>)	6-30
6.3.15.14	Register Select (<i>j_gp_regsel</i>).....	6-30
6.3.15.15	Enable Once Register Select (<i>j_en_once_regsel</i>)	6-31
6.3.15.16	External Nexus Register Select (<i>j_nexus_regsel</i>)	6-31
6.3.15.17	External Shared Nexus Control Register Select (<i>j_sncr_regsel</i>).....	6-31
6.3.15.18	External LSRL Register Select (<i>j_lsrl_regsel</i>)	6-31
6.3.15.19	Serial Data (<i>j_serial_data</i>)	6-31
6.3.15.20	Key Data In (<i>j_key_in</i>).....	6-32
6.3.16	JTAG ID Signals	6-33
6.3.16.1	JTAG ID Sequence (<i>j_id_sequence</i> [0:1])	6-33
6.3.16.2	JTAG ID Sequence (<i>j_id_sequence</i> [2:9])	6-33
6.3.16.3	JTAG ID Version (<i>j_id_version</i> [0:3])	6-34
6.4	Timing Diagrams	6-34
6.4.1	Processor Instruction/Data Transfers	6-34
6.4.1.1	Basic Read Transfer Cycles	6-35
6.4.1.2	Read Transfer with Wait State	6-36
6.4.1.3	Basic Write Transfer Cycles	6-38
6.4.1.4	Write Transfer with Wait States	6-39
6.4.1.5	Read and Write Transfers	6-41
6.4.1.6	Misaligned Accesses.....	6-45
6.4.1.7	Burst Accesses	6-48
6.4.1.8	Address Retraction.....	6-54
6.4.1.9	Error Termination Operation	6-56
6.4.2	Power Management	6-59
6.4.3	Interrupt Interface	6-60
6.4.4	Time Base Interface	6-64
6.4.5	JTAG Test Interface	6-64

Chapter 7 Power Management

7.1	Power Management	7-1
7.1.1	Active State	7-1
7.1.2	Waiting State	7-1
7.1.3	Halted State	7-1
7.1.4	Stopped State	7-2
7.1.5	Power Management Pins	7-2
7.1.6	Power Management Control Bits.....	7-2
7.1.7	Software Considerations for Power Management Using Wait Instructions	7-3
7.1.8	Software Considerations for Power Management Using Doze, Nap or Sleep	7-3

Contents

Paragraph Number	Title	Page Number
7.1.9	Debug Considerations for Power Management	7-3

Chapter 8 Debug Support

8.1	Overview	8-1
8.1.1	Software Debug Facilities	8-1
8.1.1.1	Power Architecture Book E Compatibility	8-2
8.1.2	Additional Debug Facilities	8-2
8.1.3	Hardware Debug Facilities	8-2
8.2	Software Debug Events and Exceptions	8-3
8.2.1	Instruction Address Compare Event	8-4
8.2.2	Data Address Compare Event	8-5
8.2.3	Linked Instruction Address and Data Address Compare Event	8-6
8.2.4	Trap Debug Event	8-6
8.2.5	Branch Taken Debug Event	8-6
8.2.6	Instruction Complete Debug Event	8-6
8.2.7	Interrupt Taken Debug Event	8-7
8.2.8	Critical Interrupt Taken Debug Event	8-7
8.2.9	Return Debug Event	8-7
8.2.10	Critical Return Debug Event	8-8
8.2.11	External Debug Event	8-8
8.2.12	Unconditional Debug Event	8-8
8.3	Debug Registers	8-8
8.3.1	Debug Address and Value Registers	8-8
8.3.2	Debug Control and Status Registers	8-9
8.3.2.1	Debug Control Register 0 (DBCR0)	8-9
8.3.2.2	Debug Control Register 1 (DBCR1)	8-11
8.3.2.3	Debug Control Register 2 (DBCR2)	8-13
8.3.2.4	Debug Status Register (DBSR)	8-15
8.4	External Debug Support	8-16
8.4.1	OnCE Introduction	8-17
8.4.2	JTAG/OnCE Pins	8-20
8.4.3	OnCE Internal Interface Signals	8-20
8.4.3.1	CPU Debug Request (<i>dbg_dbgrq</i>)	8-20
8.4.3.2	CPU Debug Acknowledge (<i>cpu_dbgack</i>)	8-21
8.4.3.3	CPU Address, Attributes	8-21
8.4.3.4	CPU Data	8-21
8.4.4	OnCE Interface Signals	8-21
8.4.4.1	OnCE Enable (<i>jd_en_once</i>)	8-21
8.4.4.2	OnCE Debug Request/Event (<i>jd_de_b</i> , <i>jd_de_en</i>)	8-21

Contents

Paragraph Number	Title	Page Number
8.4.4.3	e200 OnCE Debug Output (<i>jd_debug_b</i>)	8-22
8.4.4.4	e200 CPU Clock On Input (<i>jd_mclk_on</i>).....	8-22
8.4.4.5	Watchpoint Events (<i>jd_watchpt[0:5]</i>).....	8-22
8.4.5	e200 OnCE Controller and Serial Interface.....	8-23
8.4.5.1	e200 OnCE Status Register	8-23
8.4.5.2	e200 OnCE Command Register (OCMD).....	8-24
8.4.5.3	e200 OnCE Control Register (OCR).....	8-27
8.4.6	Access to Debug Resources.....	8-29
8.4.7	Methods of Entering Debug Mode	8-31
8.4.7.1	External Debug Request During RESET.....	8-31
8.4.7.2	Debug Request During RESET	8-31
8.4.7.3	Debug Request During Normal Activity	8-31
8.4.7.4	Debug Request During Waiting, Halted or Stopped State	8-31
8.4.7.5	Software Request During Normal Activity	8-32
8.4.8	CPU Status and Control Scan Chain Register (CPUSCR)	8-32
8.4.8.1	Instruction Register (IR).....	8-33
8.4.8.2	Control State Register (CTL).....	8-34
8.4.8.3	Program Counter Register (PC).....	8-37
8.4.8.4	Write-Back Bus Register (WBBRlow, WBBRhigh)	8-37
8.4.8.5	Machine State Register (MSR).....	8-38
8.4.9	Reserved Registers (Reserved)	8-38
8.5	Watchpoint Support	8-38
8.6	Basic Steps for Enabling, Using, and Exiting External Debug Mode	8-39

Chapter 9 Nexus 2+ Module

9.1	Introduction.....	9-1
9.1.1	General Description	9-1
9.1.2	Terms and Definitions.....	9-1
9.1.3	Feature List	9-2
9.1.4	Functional Block Diagram.....	9-4
9.2	Enabling Nexus 2+ Operation.....	9-4
9.3	TCODEs Supported	9-5
9.4	Nexus 2+ Programmer's Model.....	9-8
9.4.1	Client Select Control (CSC)	9-9
9.4.2	Port Configuration Register (PCR).....	9-10
9.4.3	Development Control Register 1, 2 (DC1, DC2).....	9-10
9.4.4	Development Status Register (DS).....	9-12
9.4.5	Read/Write Access Control/Status (RWCS).....	9-13
9.4.6	Read/Write Access Data (RWD)	9-14

Contents

Paragraph Number	Title	Page Number
9.4.7	Read/Write Access Address (RWA)	9-16
9.4.8	Watchpoint Trigger Register (WT)	9-16
9.5	Nexus 2+ Register Access via JTAG/OnCE	9-17
9.6	Debug Status Messages.....	9-17
9.7	Ownership Trace	9-18
9.7.1	Overview.....	9-18
9.7.2	Ownership Trace Messaging (OTM)	9-18
9.7.3	OTM Error Messages.....	9-19
9.7.4	OTM Flow	9-19
9.8	Program Trace.....	9-19
9.8.1	Branch Trace Messaging (BTM)	9-20
9.8.1.1	Zen Indirect Branch Message Instructions	9-20
9.8.1.2	Zen Direct Branch Message Instructions.....	9-21
9.8.1.3	BTM using Branch History Messages	9-21
9.8.1.4	BTM using Traditional Program Trace Messages	9-21
9.8.2	BTM Message Formats.....	9-21
9.8.2.1	Indirect Branch Messages (History)	9-22
9.8.2.2	Indirect Branch Messages (Traditional)	9-22
9.8.2.3	Direct Branch Messages (Traditional).....	9-22
9.8.2.4	Resource Full Messages	9-22
9.8.2.5	Program Correlation Messages.....	9-23
9.8.2.6	BTM Overflow Error Messages	9-24
9.8.2.7	Program Trace Synchronization Messages.....	9-24
9.8.3	BTM Operation.....	9-26
9.8.3.1	Enabling Program Trace	9-26
9.8.3.2	Relative Addressing.....	9-26
9.8.3.3	Execution Mode Indication.....	9-27
9.8.3.4	Branch/Predicate Instruction History (HIST).....	9-27
9.8.3.5	Sequential Instruction Count (I-CNT)	9-28
9.8.3.6	Program Trace Queueing	9-28
9.8.4	Program Trace Timing Diagrams (2 MDO/1 MSEO configuration).....	9-28
9.9	Watchpoint Support	9-30
9.9.1	Overview.....	9-30
9.9.2	Watchpoint Messaging.....	9-30
9.9.3	Watchpoint Error Message.....	9-31
9.9.4	Watchpoint Timing Diagram (2 MDO/1 MSEO Configuration).....	9-31
9.10	Nexus 2+ Read/Write Access to Memory-Mapped Resources.....	9-32
9.10.1	Single Write Access	9-32
9.10.2	Block Write Access.....	9-33
9.10.3	Single Read Access.....	9-33
9.10.4	Block Read Access	9-34

Contents

Paragraph Number	Title	Page Number
9.10.5	Error Handling	9-35
9.10.5.1	Bus Read/Write Error	9-35
9.10.5.2	Access Termination	9-35
9.10.6	Read/Write Access Error Message	9-35
9.11	Nexus 2+ Pin Interface	9-35
9.11.1	Pins Implemented	9-36
9.11.2	Pin Protocol.....	9-38
9.12	Rules for Output Messages	9-41
9.13	Auxiliary Port Arbitration.....	9-41
9.14	Examples.....	9-41
9.15	IEEE 1149.1 (JTAG) RD/WR Sequences	9-44
9.15.1	JTAG Sequence for Accessing Internal Nexus Registers	9-44
9.15.2	JTAG Sequence for Read Access of Memory-Mapped Resources	9-45
9.15.3	JTAG Sequence for Write Access of Memory-Mapped Resources.....	9-45

Appendix A Register Summary

Figures

Figure Number	Title	Page Number
1-1	e200z0h Block Diagram.....	1-4
1-2	e200z0 Block Diagram.....	1-5
2-1	e200z0 Supervisor Mode Programmer's Model	2-2
2-2	e200z0h Supervisor Mode Programmer's Model	2-3
2-3	e200 User Mode Program Model.....	2-4
2-4	Machine State Register (MSR)	2-7
2-5	Processor ID Register (PIR).....	2-9
2-6	Processor Version Register (PVR)	2-10
2-7	System Version Register (SVR).....	2-10
2-8	Integer Exception Register (XER)	2-11
2-9	Exception Syndrome Register (ESR).....	2-12
2-10	Machine Check Syndrome Register (MCSR).....	2-14
2-11	Hardware Implementation Dependent Register 0 (HID0)	2-16
2-12	Hardware Implementation Dependent Register 1 (HID1)	2-18
2-13	Branch Unit Control and Status Register (BUCSR)	2-19
4-1	e200z0h Block Diagram.....	4-2
4-2	e200z0 Block Diagram.....	4-3
4-3	Pipeline Diagram.....	4-6
4-4	e200 Instruction Buffers.....	4-7
4-5	e200 Instruction Buffers.....	4-7
4-6	e200 Branch Target Buffer	4-8
4-7	Basic Pipeline Flow, Single Cycle Instructions	4-9
4-8	A Load Followed By A Dependent Add Instruction	4-9
4-9	Back-to-back Load Instructions	4-10
4-10	A Load Followed By A Dependent Store Instruction.....	4-10
4-11	Basic Pipeline Flow, Branch Instructions	4-11
4-12	Basic Pipeline Flow, Branch Speculation	4-11
4-13	Basic Pipeline Flow, Multi-cycle Instructions	4-12
4-14	Pipelined Load/Store Instructions.....	4-12
4-15	Pipelined Load/Store Instructions with Wait-state.....	4-13
4-16	mtspr, mfspr Instruction Execution—(1)	4-13
4-17	mtmsr, wrtee, wrteei Instruction Execution	4-14
4-18	Interrupt Recognition and Handler Instruction Execution	4-16
4-19	Interrupt Recognition and Handler Instruction Execution— Load/Store in Progress	4-17
4-20	Interrupt Recognition and Handler Instruction Execution— Multi-Cycle Instruction Abort.....	4-18
5-1	Exception Syndrome Register (ESR).....	5-3
5-2	Machine State Register (MSR)	5-5
5-3	Machine Check Syndrome Register (MCSR).....	5-7
5-4	e200 Interrupt Vector Prefix Register (IVPR).....	5-8

Figures

Figure Number	Title	Page Number
5-5	e200 Interrupt Vector Addresses	5-9
6-1	e200z0h Signal Groups	6-3
6-2	e200z0 Signal Groups	6-4
6-3	Example External JTAG Register Design.....	6-32
6-4	Basic Read Transfers.....	6-35
6-5	Read Transfer with Wait-state.....	6-37
6-6	Basic Write Transfers	6-38
6-7	Write Transfer with Wait-State	6-40
6-8	Single Cycle Read and Write Transfers	6-41
6-9	Single Cycle Read and Write Transfers—2	6-42
6-10	Multi-Cycle Read and Write Transfers	6-43
6-11	Multi-Cycle Read and Write Transfers—2	6-44
6-12	Misaligned Read Transfer	6-45
6-13	Misaligned Write Transfer	6-46
6-14	Misaligned Write, Single Cycle Read Transfer.....	6-47
6-15	Burst Read Transfer	6-48
6-16	Burst Read with Wait-State Transfer.....	6-49
6-17	Burst Write Transfer.....	6-50
6-18	Burst Write with Wait-State Transfer	6-51
6-19	Burst Read Transfers.....	6-52
6-20	Burst Read with Wait-State Transfer, Retraction	6-53
6-21	Burst Write Transfers, Single Beat Burst.....	6-54
6-22	Read Transfer with Wait-State, Address Retraction	6-55
6-23	Burst Read with Wait-State Transfer, Retraction	6-56
6-24	Read and Write Transfers, Instr. Read Error Termination.....	6-57
6-25	Data Read Error Termination	6-58
6-26	Misaligned Write Error Termination.....	6-59
6-27	Wakeup Control Signal (<i>p_wakeup</i>)	6-59
6-28	Interrupt Interface Input Signals	6-60
6-29	e200 Interrupt Pending Operation.....	6-61
6-30	e200z0h Interrupt Pending Operation	6-61
6-31	Interrupt Acknowledge Operation	6-62
6-32	Interrupt Acknowledge Operation—2	6-63
6-33	Time Base Input Timing.....	6-64
6-34	Test Clock Input Timing	6-64
6-35	<i>j_trst_b</i> Timing.....	6-64
6-36	Test Access Port Timing	6-65
8-1	e200 Debug Resources.....	8-3
8-2	DBCR0 Register	8-9
8-3	DBCR1 Register	8-11
8-4	DBCR2 Register	8-13

Figures

Figure Number	Title	Page Number
8-5	DBSR Register	8-15
8-6	OnCE TAP Controller and Registers	8-18
8-7	e200 OnCE Controller and Serial Interface	8-23
8-8	OnCE Status Register	8-23
8-9	OnCE Command Register	8-25
8-10	OnCE Control Register	8-27
8-11	CPU Scan Chain Register (CPUSCR)	8-33
8-12	Control State Register (CTL)	8-34
9-1	Nexus 2+ Functional Block Diagram	9-4
9-2	Client Select Control Register	9-9
9-3	Port Configuration Register	9-10
9-4	Development Control Register 1	9-11
9-5	Development Control Register 2	9-12
9-6	Development Status Register	9-12
9-7	Read/Write Access Control/Status Register	9-13
9-8	Read/Write Access Data Register	9-14
9-9	Read/Write Access Address Register	9-16
9-10	Watchpoint Trigger Register	9-16
9-11	Debug Status Message Format	9-18
9-12	Ownership Trace Message Format	9-18
9-13	Error Message Format	9-19
9-14	Indirect Branch Message (History) Format	9-22
9-15	Indirect Branch Message Format	9-22
9-16	Direct Branch Message Format	9-22
9-17	Resource Full Message Format	9-23
9-18	Program Correlation Message Format	9-24
9-19	Error Message Format	9-24
9-20	Direct/Indirect Branch w/ Sync. Message Format	9-25
9-21	Indirect Branch History w/ Sync. Message Format	9-25
9-22	Relative Address Generation and Re-creation	9-27
9-23	Program Trace—Indirect Branch Message (Traditional)	9-28
9-24	Program Trace—Indirect Branch Message (History)	9-29
9-25	Program Trace—Direct Branch (Traditional) and Error Messages	9-29
9-26	Program Trace—Indirect Branch w/ Sync. Message	9-29
9-27	Watchpoint Message Format	9-30
9-28	Error Message Format	9-31
9-29	Watchpoint Message and Watchpoint Error Message	9-31
9-30	Error Message Format	9-35
9-31	Single Pin MSEO Transfers	9-39
9-32	Dual Pin MSEO Transfers	9-40
A-1	e200 User Mode Registers	A-1

Figures

Figure Number	Title	Page Number
A-2	e200z0h Supervisor Mode Registers.....	A-2
A-3	e200z0 Supervisor Mode Registers.....	A-3
A-4	Machine State Register (MSR)	A-4
A-5	Processor ID Register (PIR).....	A-4
A-6	Processor Version Register (PVR)	A-4
A-7	System Version Register (SVR).....	A-4
A-8	Integer Exception Register (XER)	A-4
A-9	Exception Syndrome Register (ESR).....	A-4
A-10	Machine Check Syndrome Register (MCSR).....	A-5
A-11	Hardware Implementation-Dependent Register 0 (HID0).....	A-5
A-12	Hardware Implementation-Dependent Register 1 (HID1).....	A-5
A-13	Branch Unit Control and Status Register (BUCSR)	A-5
A-14	DBCR0 Register	A-5
A-15	DBCR1 Register	A-5
A-16	DBCR2 Register	A-6
A-17	DBSR Register	A-6
A-18	OnCE Status Register.....	A-6
A-19	OnCE Command Register.....	A-6
A-20	OnCE Control Register	A-6
A-21	CPU Scan Chain Register (CPUSCR)	A-7
A-22	Control State Register (CTL).....	A-7
A-23	L1 Cache Configuration Register 0 (L1CFG0).....	A-7
A-24	MMU Configuration Register (MMUCFG)	A-8

Tables

Table Number	Title	Page Number
2-1	MSR Field Descriptions	2-8
2-2	PIR Field Descriptions	2-10
2-3	PVR Field Descriptions	2-10
2-4	SVR Field Descriptions	2-11
2-5	XER Field Descriptions	2-11
2-6	ESR Field Descriptions	2-12
2-7	Machine Check Syndrome Register (MCSR)	2-14
2-8	Hardware Implementation Dependent Register 0	2-16
2-9	Hardware Implementation Dependent Register 1	2-19
2-10	Branch Unit Control and Status Register	2-19
2-11	System Response to Invalid SPR Reference	2-20
2-12	Additional Synchronization Requirements for SPRs	2-21
2-13	Special Purpose Registers	2-21
2-14	Reset Settings for e200 Resources	2-23
3-1	List of Unsupported Instructions	3-1
3-2	List of Optionally Supported Instructions	3-1
3-3	List of Unimplemented SPRs	3-6
4-1	Pipeline Stages	4-5
4-2	Instruction Class Cycle Counts	4-19
4-3	Performance Effects of Storage Operand Placement	4-20
5-1	Interrupt Classifications	5-2
5-2	Exceptions and Conditions	5-2
5-3	ESR Bit Settings	5-4
5-4	MSR Bit Settings	5-5
5-5	Machine Check Syndrome Register (MCSR)	5-7
5-6	IVPR Register Fields	5-9
5-7	Hardwired Vector Offset Values	5-9
5-8	Critical Input Interrupt—Register Settings	5-10
5-9	Machine Check Interrupt—Register Settings	5-11
5-10	Data Storage Interrupt—Register Settings	5-13
5-11	Instruction Storage Interrupt—Register Settings	5-14
5-12	External Input Interrupt—Register Settings	5-14
5-13	Alignment Interrupt—Register Settings	5-15
5-14	Program Interrupt—Register Settings	5-16
5-15	System Call Interrupt—Register Settings	5-17
5-16	Debug Interrupt—Register Settings	5-19
5-17	DBSR Most Recent Reset	5-21
5-18	System Reset Interrupt—Register Settings	5-21
5-19	e200 Exception Priorities	5-23
5-20	MSR Setting Due to Interrupt	5-26
6-1	External Interface Signal Definitions	6-5

Tables

Table Number	Title	Page Number
6-2	Internal Interface Signal Definitions.....	6-8
6-3	<i>p_hrd</i> data[31:0] Byte Address Mappings.....	6-12
6-4	<i>p_d_h</i> data[31:0] Byte Address Mappings	6-12
6-5	<i>p_[d,i]_h</i> trans[1:0] Transfer Type Encoding.....	6-13
6-6	<i>p_[d,i]_h</i> size[1:0] Transfer Size Encoding.....	6-13
6-7	<i>p_[d,i]_h</i> burst[2:0] Burst Type Encoding.....	6-14
6-8	<i>p_[d,i]_h</i> prot[5:0] Protection Control Encoding	6-14
6-9	<i>p_hb</i> strb[3:0] to Byte Address Mappings.....	6-15
6-10	Byte Strobe Assertion for Transfers.....	6-16
6-11	Big-Endian Memory Storage	6-17
6-12	<i>p_d_h</i> resp[2:0] Transfer Response Encoding	6-18
6-13	<i>p_i_h</i> resp[1:0] Transfer Response Encoding.....	6-19
6-14	Processor Status Encoding	6-22
6-15	e200 Debug/Emulation Support Signals	6-26
6-16	e200 Development Support (Nexus2+) Signals.....	6-28
6-17	JTAG Primary Interface Signals	6-28
6-18	JTAG Signals Used to Support External Registers	6-29
6-19	JTAG General Purpose Register Select Decoding.....	6-30
6-20	JTAG Register ID Fields.....	6-33
6-21	JTAG ID Register Inputs.....	6-33
8-1	DBCR0 Bit Definitions.....	8-10
8-2	DBCR1 Bit Definitions.....	8-12
8-3	DBCR2 Bit Definitions.....	8-14
8-4	DBSR Bit Definitions	8-15
8-5	JTAG/OnCE Primary Interface Signals	8-20
8-6	OnCE Status Register Bit Definitions.....	8-24
8-7	OnCE Command Register Bit Definitions.....	8-25
8-8	e200 OnCE Register Addressing	8-26
8-9	OnCE Control Register Bit Definitions	8-28
8-10	OnCE Register Access Requirements.....	8-30
8-11	Watchpoint Output Signal Assignments	8-38
9-1	Terms and Definitions.....	9-1
9-2	Public TCODEs Supported	9-5
9-3	Error Code Encoding (TCODE = 8)	9-7
9-4	RCODE values (TCODE = 27).....	9-7
9-5	Event Code Encoding (TCODE = 33)	9-7
9-6	Data Trace Size Encodings (TCODE = 5,6,13,14)	9-8
9-7	Nexus 2+ Register Map.....	9-9
9-8	Client Select Control Register Fields.....	9-9
9-9	Port Configuration Register Fields	9-10
9-10	Development Control Register 1 Fields.....	9-11

Tables

Table Number	Title	Page Number
9-11	Development Control Register 2 Fields.....	9-12
9-12	Development Status Register Fields	9-13
9-13	Read/Write Access Control/Status Register Fields.....	9-13
9-14	Read/Write Access Status Bit Encoding	9-14
9-15	RWD data placement for Transfers.....	9-15
9-16	RWD byte lane data placement.....	9-15
9-17	Watchpoint Trigger Register Fields	9-16
9-18	Indirect Branch Message Sources	9-20
9-19	Direct Branch Message Sources	9-21
9-20	RCODE Encoding.....	9-23
9-21	Program Trace Exception Summary	9-25
9-22	Watchpoint Source Encoding.....	9-30
9-23	JTAG Pins for Nexus 2+	9-36
9-24	Nexus 2+ Auxiliary Pins	9-37
9-25	Nexus Port Arbitration Signals	9-38
9-26	MSEO Pin(s) Protocol	9-38
9-27	MDO Request Encodings.....	9-41
9-28	Indirect Branch Message Example (2 MDO/1 MSEO)	9-42
9-29	Indirect Branch Message Example (8 MDO/2 MSEO)	9-42
9-30	Direct Branch Message Example (2 MDO/1 MSEO).....	9-43
9-31	Direct Branch Message Example (8 MDO/2 MSEO).....	9-43
9-32	Data Write Message Example (8 MDO/1 MSEO).....	9-43
9-33	Data Write Message Example (8 MDO/2 MSEO).....	9-44
9-34	Accessing Internal Nexus 2+ Registers via JTAG/OnCE	9-44
9-35	Accessing Memory-Mapped Resources (Reads)	9-45
9-36	Accessing Memory-Mapped Resources (Writes)	9-45



Tables

Table Number	Title	Page Number
-----------------	-------	----------------

Chapter 1

e200z0 and e200z0h Overview

1.1 Overview of the e200z0 and e200z0h Cores

The e200 processor family is a set of CPU cores that implement low-cost versions of the Power Architecture™ Book E architecture. e200 processors are designed for deeply embedded control applications, which require low cost solutions rather than maximum performance.

The e200z0 and e200z0h processors integrate an integer execution unit, branch control unit, instruction fetch and load/store units, and a multi-ported register file capable of sustaining three read and two write operations per clock. Most integer instructions execute in a single clock cycle. Branch target prefetching is performed by the branch unit to allow single-cycle branches in some cases.

The e200z0 and e200z0h cores are single-issue, 32-bit Power Architecture Book EVLE-only designs with 32-bit general purpose registers (GPRs). All arithmetic instructions that execute in the core operate on data in the general purpose registers (GPRs).

Instead of the base Power Architecture Book E instruction set support, the e200z0 and e200z0h cores only implement the VLE (variable-length encoding) APU, providing improved code density. The VLE APU is further documented in the *PowerPC™ VLE APU Definition, Version 1.01*, a separate document.

In the remainder of this document, the e200z0 and e200z0h core are also referred to as the “e200z0 core” or “e200 core” when referring to the whole e200 family. The term ‘e200z0h’ is used where differences exist between the e200z0 and e200z0h.

1.1.1 Features

The following is a list of some of the key features of the e200z0 and e200z0h cores:

- 32-bit Power Architecture Book EVLE-only programmer’s model
- Single issue, 32-bit CPU
- Implements the VLE APU for reduced code footprint
- In-order execution and retirement
- Precise exception handling
- Branch processing unit
 - Dedicated branch address calculation adder
 - Branch acceleration using Branch Target Buffer (e200z0h only)
- Supports independent instruction and data accesses to different memory subsystems, such as SRAM and Flash memory via independent Instruction and Data bus interface units (BIUs) (e200z0h only).

- Supports instruction and data access via a unified 32-bit Instruction/Data BIU (e200z0 only).
- Load/store unit
 - 1 cycle load latency
 - Fully pipelined
 - Big-endian support only
 - Misaligned access support
 - Zero load-to-use pipeline bubbles for aligned transfers
- Power management
 - Low power design
 - Power saving modes: doze, nap, sleep, and wait
 - Dynamic power management of execution units
- Testability
 - Synthesizeable, full MuxD scan design
 - ABIST/MBIST for optional memory arrays

1.1.2 Microarchitecture Summary

The e200 processor utilizes a four stage pipeline for instruction execution. The Instruction Fetch (stage 1), Instruction Decode/Register file Read/Effective Address Calculation (stage 2), Execute/Memory Access (stage 3), and Register Writeback (stage 4) stages operate in an overlapped fashion, allowing single clock instruction execution for most instructions.

The integer execution unit consists of a 32-bit Arithmetic Unit (AU), a Logic Unit (LU), a 32-bit Barrel shifter (Shifter), a Mask-Insertion Unit (MIU), a Condition Register manipulation Unit (CRU), a Count-Leading-Zeros unit (CLZ), an 8x32 Hardware Multiplier array, result feed-forward hardware, and a hardware divider.

Arithmetic and logical operations are executed in a single cycle with the exception of the divide and multiply instructions. A Count-Leading-Zeros unit operates in a single clock cycle.

The Instruction Unit contains a PC incrementer and a dedicated Branch Address adder to minimize delays during change of flow operations. Sequential prefetching is performed to ensure a supply of instructions into the execution pipeline. Branch target prefetching from the BTB is performed to accelerate certain taken branches in the e200z30h. Prefetched instructions are placed into an instruction buffer with 2 entries in e200z0 and 4 entries in e200z0h, each capable of holding a single 32-bit instruction or a pair of 16-bit instructions.

Conditional branches which are not taken execute in a single clock. Branches with successful target prefetching have an effective execution time of 1 clock on e200z0h. All other taken branches have an execution time of two clocks.

Memory load and store operations are provided for byte, halfword, and word (32-bit) data with automatic zero or sign extension of byte and halfword load data as well as optional byte reversal of data. These instructions can be pipelined to allow effective single cycle throughput. Load and store multiple word instructions allow low overhead context save and restore operations. The load/store unit contains a

dedicated effective address adder to allow effective address generation to be optimized. Also, a load-to-use dependency does not incur any pipeline bubbles for most cases.

The Condition Register unit supports the condition register (CR) and condition register operations defined by the PowerPC™ architecture. The condition register consists of eight 4-bit fields that reflect the results of certain operations, such as move, integer and floating-point compare, arithmetic, and logical instructions, and provide a mechanism for testing and branching.

Vectored and autovectored interrupts are supported by the CPU. Vectored interrupt support is provided to allow multiple interrupt sources to have unique interrupt handlers invoked with no software overhead.




Figure 1-1. e200z0h Block Diagram




Figure 1-2. e200z0 Block Diagram

1.1.2.1 Instruction Unit Features

The features of the e200 Instruction unit are:

- 32-bit instruction fetch path supports fetching of one 32-bit instruction per clock, or up to two 16-bit VLE instructions per clock.

- Instruction buffer with 2 entries in e200z0 and 4 entries in e200z0h, each holding a single 32-bit instruction, or a pair of 16-bit instructions
- Dedicated PC incrementer supporting instruction prefetches
- Branch unit with dedicated branch address adder supporting single cycle of execution of certain branches, two cycles for all others

1.1.2.2 Integer Unit Features

The e200 integer unit supports single cycle execution of most integer instructions:

- 32-bit AU for arithmetic and comparison operations
- 32-bit LU for logical operations
- 32-bit priority encoder for count leading zero's function
- 32-bit single cycle barrel shifter for shifts and rotates
- 32-bit mask unit for data masking and insertion
- Divider logic for signed and unsigned divide in 5 to 34 clocks with minimized execution timing
- 8x32 hardware multiplier array supports 1 to 4 cycle 32x32->32 multiply (early out)

1.1.2.3 Load/Store Unit Features

The e200 load/store unit supports load, store, and the load multiple / store multiple instructions:

- 32-bit effective address adder for data memory address calculations
- Pipelined operation supports throughput of one load or store operation per cycle
- 32-bit interface to memory dedicated memory interface on e200z0h)

1.1.2.4 e200z0h System Bus Features

The features of the e200z30h System Bus interface are as follows:

- Independent Instruction and Data Buses
- AMBA AHB Lite Rev 2.0 Specification with support for ARM v6 AMBA Extensions
 - Exclusive Access Monitor
 - Byte Lane Strobes
 - Cache Allocate Support
- 32-bit address bus plus attributes and control on each bus
- 32-bit read data bus for Instruction Interface
- Separate uni-directional 32-bit read data bus and 32-bit write data bus for Data Interface
- Overlapped, in-order accesses

1.1.2.5 e200z0 System Bus Features

The features of the e200z0 System Bus interface are as follows:

- Unified Instruction/Data Bus

- AMBA AHB2.v6 protocol
- 32-bit address bus plus attributes and control
- Separate uni-directional 32-bit read data bus and 32-bit write data bus
- Overlapped, in-order accesses

1.1.2.6 Nexus 2+ Features

The module is compatible with Class 2 of the IEEE-ISTO Std 5001™-2003, with additional Class 3 and Class 4 features available. The following features are implemented:

- Program Trace via Branch Trace Messaging (BTM). Branch trace messaging displays program flow discontinuities (direct and indirect branches, exceptions, etc.), allowing the development tool to interpolate what transpires between the discontinuities. Thus, static code may be traced.
- Ownership Trace via Ownership Trace Messaging (OTM). OTM facilitates ownership trace by providing visibility of which process ID or operating system task is activated. An Ownership Trace Message is transmitted when a new process/task is activated, allowing the development tool to trace ownership flow.
- Run-time access to the processor memory map via the JTAG port. This allows for enhanced download/upload capabilities.
- Watchpoint Messaging via the auxiliary interface
- Watchpoint Trigger enable of Program Trace Messaging
- Auxiliary interface for higher data input/output
 - Configurable (min/max) Message Data Out pins (*nex_mdo[n:0]*)
 - One (1) or two (2) Message Start/End Out pins (*nex_mseo_b[1:0]*)
 - One (1) Read/Write Ready pin (*nex_rdy_b*) pin
 - One (1) Watchpoint Event pin (*nex_evto_b*)
 - One (1) Event In pin (*nex_evti_b*)
 - One (1) MCKO (Message Clock Out) pin
- Registers for Program Trace, Ownership Trace and Watchpoint Trigger control.
- All features controllable and configurable via the JTAG port



Chapter 2

Register Model

This section describes the registers implemented in the e200z0 and e200z0h cores. It includes an overview of registers defined by the PowerPC Book E architecture, highlighting differences in how these registers are implemented in the e200 core, and provides a detailed description of e200-specific registers. Full descriptions of the architecture-defined register set are provided in Power Architecture Book E Specification.

The Power Architecture Book E defines register-to-register operations for all computational instructions. Source data for these instructions are accessed from the on-chip registers or are provided as immediate values embedded in the opcode. The three-register instruction format allows specification of a target register distinct from the two source registers, thus preserving the original data for use by other instructions. Data is transferred between memory and registers with explicit load and store instructions only.

[Figure 2-1](#), [Figure 2-2](#), and [Figure 2-3](#) show the e200 register set including the registers which are accessible while in supervisor mode, and the registers which are accessible in user mode. The number to the right of the special-purpose registers (SPRs) is the decimal number used in the instruction syntax to access the register (for example, the integer exception register (XER) is SPR 1).

NOTE

e200z0 and e200z0h are a 32-bit implementation of the Power Architecture Book E specification. In this document, register bits are sometimes numbered from bit 0 (most significant bit) to 31 (least significant bit), rather than the Book E numbering scheme of 32:63, thus register bit numbers for some registers in Book E are 32 higher.

Where appropriate, the Book E defined bit numbers are shown in parentheses.





Figure 2-1. e200z0 Supervisor Mode Programmer's Model



Processor Control Registers

Machine State

MSR

Processor Version

PVR

SPR 287

Processor ID

PIR

SPR 286

System Version¹

SVR

SPR 1023

Hardware Implementation Dependent¹

HID0

SPR 1008

HID1

SPR 1009

Debug Registers²

Debug Control

DBCR0

SPR 308

DBCR1

SPR 309

DBCR2

SPR 310

Debug Status

DBSR

SPR 304

Instruction Address Compare

IAC1

SPR 312

IAC2

SPR 313

IAC3

SPR 314

IAC4

SPR 315

Data Address Compare

DAC1

SPR 316

DAC2

SPR 317

Memory Management Registers

Process ID

PID0

SPR 48

Configuration (Read-only)

MMUCFG

SPR 1015

Cache Registers

Cache Configuration (Read-only)

L1CFG0

SPR 515

BTB Registers

BTB Control¹

BUCSR

SPR 1013

1 - These e200-specific registers may not be supported by other Power Architecture processors

2 - Optional registers defined by the Power Architecture Book E

Figure 2-2. e200z0h Supervisor Mode Programmer's Model




Figure 2-3. e200 User Mode Program Model

General purpose registers (GPRs) are accessed through instruction operands. Access to other registers can be explicit (by using instructions for that purpose such as Move to Special Purpose Register (**mtspr**) and Move from Special Purpose Register (**mfspir**) instructions) or implicit as part of the execution of an instruction. Some registers are accessed both explicitly and implicitly.

2.1 Power Architecture Book E Registers

e200 supports a subset of the registers defined by *Power Architecture™ Book E Specification*. Notable exceptions are the Floating Point registers FPR0–FPR31 and FPSCR. e200z0 and e200z0h cores do not support the Book E floating-point architecture. The e200-supported Power Architecture Book E registers are described as follows (e200-specific registers are described in the [Section 2.2, “e200-Specific Special Purpose Registers”](#)).

- User-level registers—The user-level registers can be accessed by all software with either user or supervisor privileges. They include the following:
 - General-purpose registers (GPRs). The thirty-two 32-bit GPRs (GPR0–GPR31) serve as data source or destination registers for integer instructions and provide data for generating addresses.
 - Condition register (CR). The 32-bit CR consists of eight 4-bit fields, CR0–CR7, that reflect results of certain arithmetic operations and provide a mechanism for testing and branching. See “Condition Register (CR),” in Chapter 3, “Branch and Condition Register Operations, *Power Architecture Book E Specification*.”

The remaining user-level registers are SPRs. Note that the Power Architecture Book E provides the **mtspr** and **mfspir** instructions for accessing SPRs.

- Integer exception register (XER). The XER indicates overflow and carries for integer operations. See “XER Register (XER),” in Chapter 4, “Integer Operations” of *Power Architecture Book E Specification* for more information.

- Link register (LR). The LR provides the branch target address for the Branch to Link Register (**se_blr**, **se_bldr**) instructions, and is used to hold the address of the instruction that follows a branch and link instruction, typically used for linking to subroutines. See “Link Register (LR)”, in Chapter 3, “Branch and Condition Register Operations” of *Power Architecture Book E Specification*.
- Count register (CTR). The CTR holds a loop count that can be decremented during execution of appropriately coded branch instructions. The CTR also provides the branch target address for the Branch to Count Register (**se_bctr**, **se_bctrl**) instructions. See “Count Register (CTR)”, in Chapter 3, “Branch and Condition Register Operations” of *Power Architecture Book E Specification*.
- Supervisor-level registers—In addition to the registers accessible in user mode, Supervisor-level software has access to additional control and status registers used for configuration, exception handling, and other operating system functions. The Power Architecture Book E defines the following supervisor-level registers:
 - Processor Control registers
 - Machine State Register (MSR). The MSR defines the state of the processor. The MSR can be modified by the Move to Machine State Register (**mtmsr**), System Call (**se_sc**), and Return from Exception (**se_rfi**, **se_rfci**, **se_rfdi**) instructions. It can be read by the Move from Machine State Register (**mfmsr**) instruction. When an interrupt occurs, the contents of the MSR are saved to one of the machine state save/restore registers (SRR1, CSRR1, DSRR1).
 - Processor version register (PVR). This register is a read-only register that identifies the version (model) and revision level of the processor.
 - Processor Identification Register (PIR). This read-only register is provided to distinguish the processor from other processors in the system.
 - Storage Control register
 - Process ID Register (PID, also referred to as PID0). This register is provided to indicate the current process or task identifier. It is used by the optional MMU as an extension to the effective address, and by the optional external Nexus2/3 modules for Ownership Trace message generation. Although the Power Architecture Book E allows for multiple PIDs, the e200z0 and e200z0h implement only one.
 - Interrupt Registers
 - Data Exception Address Register (DEAR). After most Data Storage Interrupts (DSI), or on an Alignment Interrupt, the DEAR is set to the effective address (EA) generated by the faulting instruction.
 - SPRG0-SPRG1. The SPRG0-SPRG1 registers are provided for operating system or interrupt handler use.
 - Exception Syndrome Register (ESR). The ESR register provides a syndrome to differentiate between the different kinds of exceptions which can generate the same interrupt.
 - Interrupt Vector Prefix Register (IVPR). This register together with hardwired offsets which replace the IVOR0-15 registers provide the address of the interrupt handler for different classes of interrupts.

- Save/Restore Register 0 (SRR0). The SRR0 register is used to save machine state on a non-critical interrupt, and contains the address of the instruction at which execution resumes when an **se_rfi** instruction is executed at the end of a non-critical class interrupt handler routine.
- Critical Save/Restore register 0 (CSRR0). The CSRR0 register is used to save machine state on a critical interrupt, and contains the address of the instruction at which execution resumes when an **se_rfci** instruction is executed at the end of a critical class interrupt handler routine.
- Save/Restore register 1 (SRR1). The SRR1 register is used to save machine state from the MSR on non-critical interrupts, and to restore machine state when **se_rfi** executes.
- Critical Save/Restore register 1 (CSRR1). The CSRR1 register is used to save machine state from the MSR on critical interrupts, and to restore machine state when **se_rfci** executes.
- Debug facility registers
 - Debug Control Registers (DBCR0-DBCR2). These registers provide control for enabling and configuring debug events.
 - Debug Status Register (DBSR). This register contains debug event status.
 - Instruction Address Compare registers (IAC1-IAC4). These registers contain addresses and/or masks which are used to specify Instruction Address Compare debug events.
 - Data address compare registers (DAC1-2). These registers contain addresses and/or masks which are used to specify Data Address Compare debug events.
 - e200 does **not** implement the Data Value Compare registers (DVC1 and DVC2).

2.2 e200-Specific Special Purpose Registers

The PowerPC Book E architecture allows implementation-specific special purpose registers. Those incorporated in the e200 core are as follows:

- User-level registers—The user-level registers can be accessed by all software with either user or supervisor privileges. They include the following:
 - The L1 Cache Configuration register (L1CFG0). This read-only register allows software to query the configuration of the L1 Cache. For the e200z0 and e200z0h, this register returns all zeros indicating no cache is present.
- Supervisor-level registers—The following supervisor-level registers are defined in e200 in addition to the Power Architecture Book E registers described above:
 - Configuration Registers
 - Hardware implementation-dependent register 0 (HID0). This register controls various processor and system functions.
 - Hardware implementation-dependent register 1 (HID1). This register controls various processor and system functions.
 - Exception Handling and Control Registers
 - Machine Check Syndrome register (MCSR). This register provides a syndrome to differentiate between the different kinds of conditions which can generate a Machine Check.

- Debug Save/Restore register 0 (DSRR0). When enabled, the DSRR0 register is used to save the address of the instruction at which execution continues when **se_rfdi** executes at the end of a debug interrupt handler routine.
- Debug Save/Restore register 1 (DSRR1). When enabled, the DSRR1 register is used to save machine status on debug interrupts and to restore machine status when **se_rfdi** executes.
- Branch Unit Control and Status Register (BUCSR) controls operation of the BTB in e200z0 and e200z0h.
- L1 Cache Configuration Register (L1CFG0) is a read-only register that allows software to query the configuration of the L1 Cache. For the e200z0 and e200z0h this register returns all zeros.
- MMU Configuration Register (MMUCFG) is a read-only register that allows software to query the configuration of the MMU. For e200z0 and e200z0h, this register returns the value 0x0000_0003 indicating no MMU is present.
- System version register (SVR). This register is a read-only register that identifies the version (model) and revision level of the SoC which includes an e200 Power Architecture processor.

Note that it is not guaranteed that the implementation of e200 core-specific registers is consistent among Power Architecture processors, although other processors may implement similar or identical registers. All e200 SPR definitions are compliant with the Freescale EIS specification definitions as documented in the *EREF: A Programmer's Reference Manual for Freescale Book E Processors*.

2.3 Special Purpose Register Descriptions

2.3.1 Machine State Register (MSR)

The Machine State Register defines the state of the processor. [Chapter 5, “Interrupts and Exceptions,”](#) describes how the MSR is affected when Interrupts occur. The e200 MSR is shown in [Figure 2-4](#).

e200z0h	0	UCLE	Allocated	0	WE	CE	0	EE	PR	FP	ME	FE0	0	DE	FE1	0	IS	DS	0	RI	0
e200z0	0	UCLE	Allocated	0	WE	CE	0	EE	PR	FP	ME	FE0	0	DE	FE1	0	IS	DS	0		
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31																					
																					Read/Write; Reset—0x0

Figure 2-4. Machine State Register (MSR)

The MSR bits are defined in [Table 2-1](#).

Table 2-1. MSR Field Descriptions

Bit(s)	Name	Description
0:4 (32:36)	—	Reserved ¹
5 (37)	UCLE ²	User Cache Lock Enable 0 Execution of the cache locking instructions in user mode (MSR[PR]=1) disabled; DSI exception taken instead, and ILK or DLK set in ESR. 1 Execution of the cache lock instructions in user mode enabled.
6 (38)	Allocated	Allocated ³ - Allocated for SPE Not supported on e200z0 or e200z0h
7:12 (39:44)	—	Reserved ¹
13 (45)	WE	Wait State (Power management) enable 0 Power management is disabled. 1 Power management is enabled. The processor can enter a power-saving mode when additional conditions are present. The mode chosen is determined by the DOZE, NAP, and SLEEP bits in the HID0 register, described in Section 2.3.9, “Hardware Implementation Dependent Register 0 (HID0).”
14 (46)	CE	Critical Interrupt Enable 0 Critical Input interrupts are disabled. 1 Critical Input interrupts are enabled.
15 (47)	—	Reserved ¹
16 (48)	EE	External Interrupt Enable 0 External Input interrupts are disabled. 1 External Input interrupts are enabled.
17 (49)	PR	Problem State 0 The processor is in supervisor mode, can execute any instruction, and can access any resource (for example, GPRs, SPRs, MSR, etc.). 1 The processor is in user mode, cannot execute any privileged instruction, and cannot access any privileged resource.
18 (50)	FP ⁴	Floating-Point Available 0 Floating point unit is unavailable. The processor cannot execute floating-point instructions, including floating-point loads, stores, and moves. (A FP Unavailable interrupt is generated on attempted execution of floating point instructions). 1 Floating Point unit is available. The processor can execute floating-point instructions.
19 (51)	ME	Machine Check Enable 0 Machine Check interrupts are disabled. 1 Machine Check interrupts are enabled.
20 (52)	FE0	Floating-point exception mode 0 (not used by e200)
21 (53)	—	Reserved ¹
22 (54)	DE	Debug Interrupt Enable 0 Debug interrupts are disabled. 1 Debug interrupts are enabled.

Table 2-1. MSR Field Descriptions (continued)

Bit(s)	Name	Description
23 (55)	FE1	Floating-point exception mode 1 (not used by e200)
24 (56)	—	Reserved ¹
25 (57)	—	Reserved ¹
26 (58)	IS	Instruction Address Space 0 The processor directs all instruction fetches to address space 0. 1 The processor directs all instruction fetches to address space 1.
27 (59)	DS	Data Address Space 0 The processor directs all data storage accesses to address space 0. 1 The processor directs all data storage accesses to address space 1.
28:29 (60:61)	—	Reserved ¹
30 (62)	RI	Recoverable Interrupt 0 Machine Check interrupt is not recoverable. 1 Machine Check interrupt may be recoverable. This bit is cleared when a Machine check interrupt is taken, or when a critical class interrupt using CSRR0/1 is taken. It is not set by hardware, and does not affect processor operation. It is provided as a software assist.
30:31 (62:63)	—	Reserved ¹

¹ This bit is not implemented, is read as zero, and writes are ignored.

² This bit is implemented but ignored because no cache is implemented

³ This bits is should be written with zero for future compatibility.

⁴ This bit is implemented but ignored

2.3.2 Processor ID Register (PIR)

The processor ID for the processor core is contained in the Processor ID Register (PIR), shown in [Figure 2-5](#). The contents of the PIR register are a reflection of hardware input signals to the e200 core. This register is read-only.




Figure 2-5. Processor ID Register (PIR)

The PIR fields are defined in [Table 2-2](#).

Table 2-2. PIR Field Descriptions

Bits	Name	Description
0:23	—	These bits always reads 0.
24:31	ID	These bits are a reflection of the values provided on the <i>p_cpuid</i> [0:7] input signals.

2.3.3 Processor Version Register (PVR)

The Processor Version Register (PVR), shown in [Figure 2-6](#), contains the processor version number for the processor core.




Figure 2-6. Processor Version Register (PVR)

The PVR bit fields are shown in [Table 2-3](#).

Table 2-3. PVR Field Descriptions

Bits	Name	Description
0–5	Reserved	Reserved
6–11	Type	A 6-bit number that, together with the version number, uniquely identifies a particular processor version.
12–15	Version	A 4-bit number that, together with the version number, uniquely identifies a particular processor version.
16–23	Reserved	Reserved
24–27	Major Rev	A 4-bit number that, together with the ID number, distinguishes between various releases of a particular version (that is, an engineering change level). The value of the revision portion of the PVR is implementation-specific. The processor revision level is changed for each revision of the device.
28–31	MBG ID	A 4-bit number that, together with the major revision number, distinguishes between various releases of a particular version (that is, an engineering change level). The value of the revision portion of the PVR is implementation-specific. The processor revision level is changed for each revision of the device.

2.3.4 System Version Register (SVR)

The System Version Register (SVR), shown in [Figure 2-7](#), contains system version information for an e200-based SoC.




Figure 2-7. System Version Register (SVR)

This register is used to specify a particular implementation of an e200-based system. This register is read-only. The SVR bit fields are shown in [Table 2-4](#).

Table 2-4. SVR Field Descriptions

Bits	Name	Description
0–31	Version	SVR number is SoC specific.

2.3.5 Integer Exception Register (XER)

The XER bit assignments are shown in [Figure 2-8](#).




Figure 2-8. Integer Exception Register (XER)

The XER fields are defined in [Table 2-5](#).

Table 2-5. XER Field Descriptions

Bits	Name	Description
0 (32)	SO	Summary Overflow (per Book E)
1 (33)	OV	Overflow (per Book E)
2 (34)	CA	Carry (per Book E)
3:24 (35:56)	—	Reserved ¹
25:31 (57:63)	Bytecnt ²	Preserved for lswi , lswx , stswi , stswx string instructions

¹ These bits are not implemented, is read as zero, and writes are ignored.

² These bits are implemented to support emulation of the string instructions.

2.3.6 Exception Syndrome Register

The Exception Syndrome Register (ESR) provides a syndrome to differentiate between exceptions that can generate the same interrupt type. e200 adds some implementation-specific bits to this register, as seen in [Figure 2-9](#).




Figure 2-9. Exception Syndrome Register (ESR)

The ESR fields are defined in [Table 2-6](#).

Table 2-6. ESR Field Descriptions

Bit(s)	Name	Description	Associated Interrupt Type
0:3 (32:35)	—	Allocated ¹	—
4 (36)	PIL	Illegal Instruction exception	Program
5 (37)	PPR	Privileged Instruction exception	Program
6 (38)	PTR	Trap exception	Program
7 (39)	FP ²	Floating-point operation	Alignment Data Storage Data TLB Program
8 (40)	ST	Store operation	Alignment Data Storage Data TLB
9 (41)	—	Reserved ¹	—
10 (42)	DLK ²	Data Cache Locking	Data Storage
11 (43)	ILK ²	Instruction Cache Locking	Data Storage
12 (44)	AP	Auxiliary Processor operation (Currently unused in e200)	Alignment Data Storage Data TLB Program
13 (45)	PUO	Unimplemented Operation exception	Program

Table 2-6. ESR Field Descriptions (continued)

Bit(s)	Name	Description	Associated Interrupt Type
14 (46)	BO	Byte Ordering exception Mismatched Instruction Storage exception	Data Storage Instruction Storage
15 (47)	PIE	Program Imprecise exception (Reserved)	Currently unused in e200
16:23 (48:55)	—	Reserved ¹	—
24 (56)	EFP	Embedded Floating-point APU Operation	Allocated, not set by hardware
25 (57)	—	Allocated ¹	—
26 (58)	VLEMI	VLE Mode Instruction	Data Storage Instruction Storage Alignment Program System Call
27:29 (59:61)	—	Allocated ¹	—
30 (62)	MIF ²	Misaligned Instruction Fetch	Instruction Storage Instruction TLB
31 (63)	XTE	External Termination Error (Precise)	Data Storage Instruction Storage

¹ These bits are not implemented and should be written with zero for future compatibility.

² Unused on e200z0h and e200z0.

2.3.6.1 Power Architecture VLE Mode Instruction Syndrome

The ESR[VLEMI] bit is provided to indicate that an interrupt was caused by a Power Architecture VLE instruction. This syndrome bit is set on an exception associated with execution or attempted execution of a Power Architecture VLE instruction. This bit is updated for the interrupt types indicated in [Table 2-6](#).

2.3.6.2 Misaligned Instruction Fetch Syndrome

The ESR[MIF] bit is provided to indicate that an Instruction Storage Interrupt was caused by an attempt to fetch an instruction from a Book E page which was not aligned on a word boundary. The fetch may have been caused by execution of a Branch class instruction from a VLE page to a non-VLE page, a Branch to LR instruction with LR[62]=1, a Branch to CTR instruction with CTR[62]=1, execution of an **se_rfi** instruction with SRR0[62]=1, execution of an **se_rfci** instruction with CSRR0[62]=1, or execution of an **se_rfdi** instruction with DSRR0[62]=1, where the destination address corresponds to an instruction page which is not marked as a Power Architecture VLE page.

The ESR[MIF] bit is also used to indicate that an Instruction TLB Interrupt was caused by a TLB miss on the second half of a misaligned 32-bit Power Architecture VLE Instruction. For this case, SRR0 points to the first half of the instruction, which resides on the previous page from the miss at page offset 0xFFE. The

ITLB handler may need to realize that the miss corresponds to the next page, although MMU MAS2 contents correctly reflects the page corresponding to the miss.

NOTE

This bit is allocated, but is not set by e200z0 and e200z0h because no Book E pages exist and no MMU is implemented on these cores.

2.3.6.3 Precise External Termination Error Syndrome

The ESR[XTE] bit is provided to indicate that a precise external termination error DSI or ISI interrupt was caused by an instruction. This syndrome bit is set on an external termination error exception that is reported in a precise manner via a DSI or ISI as opposed to a machine check.

2.3.7 Machine Check Syndrome Register (MCSR)

When the core complex takes a machine check interrupt, it updates the Machine Check Syndrome register (MCSR) to differentiate between machine check conditions. The MCSR is shown in [Figure 2-10](#).




Figure 2-10. Machine Check Syndrome Register (MCSR)

[Table 2-7](#) describes MCSR fields. The MCSR indicates the source of a machine check condition is recoverable. When a syndrome bit in the MCSR is set, the core complex asserts *p_mcp_out* for system information.

Table 2-7. Machine Check Syndrome Register (MCSR)

Bit	Name	Description
0 (32)	MCP	Machine check input pin
1 (33)	—	Reserved, should be cleared.
2 (34)	CP_PERR ¹	Cache push parity error

Table 2-7. Machine Check Syndrome Register (MCSR) (continued)

Bit	Name	Description
3 (35)	CPERR ¹	Cache parity error
4 (36)	EXCP_ERR	Bus Error on first instruction fetch for an exception handler
5:26 (37:58)	—	Reserved. Should be cleared.
e200z0h: 5:10 (37:42) e200z0: 5:26 (37:58)	—	Reserved, should be cleared.
11 (43)	NMI	Non-maskable interrupt input pin
12:26 (44:58)	—	Reserved, should be cleared.
27 (59)	BUS_IRERR	Read bus error on Instruction fetch
28 (60)	BUS_DRERR	Read bus error on data load
29 (61)	BUS_WRERR	Write bus error on buffered store
30:31 (62:63)	—	Reserved, should be cleared.

¹ Unused on e200z0 and e200z0h

2.3.8 Debug Registers

The Debug facility registers are described in [Chapter 8, “Debug Support.”](#)

2.3.9 Hardware Implementation Dependent Register 0 (HID0)

The HID0 register is an e200 implementation dependent register used for various configuration and control functions. The HID0 register is shown in [Figure 2-11](#).




Figure 2-11. Hardware Implementation Dependent Register 0 (HID0)

The HID0 fields are defined in [Table 2-8](#).

Table 2-8. Hardware Implementation Dependent Register 0

Bits	Name	Description
0	EMCP	Enable machine check pin (<i>p_mcp_b</i>) 0 <i>p_mcp_b</i> pin is disabled. 1 <i>p_mcp_b</i> pin is enabled. If MSR[ME] = 0, asserting <i>p_mcp_b</i> causes checkstop. If MSR[ME] = 1, asserting <i>p_mcp_b</i> causes a machine check interrupt. The primary purpose of this bit is to mask out further machine check exceptions caused by assertion of <i>p_mcp_b</i> .
1:5	—	Reserved ¹
6:7	BPRED ²	Branch Prediction (Acceleration) Control 00 - Branch acceleration is enabled. 01 - Branch acceleration is disabled for backward branches. 10 - Branch acceleration is disabled for forward branches. 11 - Branch acceleration is disabled for both branch directions. This field controls instruction buffer lookahead for branch acceleration. Note that for branches with "AA" = '1', the MSB of the displacement field is still used to indicate forward/backward, even though the branch is absolute. This field is used in conjunction with the BUCSR.
8	DOZE	Configure for Doze power management mode 0 Doze mode is disabled 1 Doze mode is enabled Doze mode is invoked by setting MSR[WE] while this bit is set.
9	NAP	Configure for Nap power management mode 0 Nap mode is disabled 1 Nap mode is enabled Nap mode is invoked by setting MSR[WE] while this bit is set.

Table 2-8. Hardware Implementation Dependent Register 0 (continued)

Bits	Name	Description
10	SLEEP	Configure for Sleep power management mode 0 Sleep mode is disabled 1 Sleep mode is enabled Sleep mode is invoked by setting MSR[WE] while this bit is set. Only one of DOZE, NAP, or SLEEP should be set for proper operation.
11:13	—	Reserved ¹
14	ICR	Interrupt Inputs Clear Reservation 0 External Input, Critical Input, and Non-Maskable Interrupts do not affect reservation status 1 External Input, Critical Input, and Non-Maskable Interrupts clear an outstanding reservation
15	NHR	Not hardware reset 0 Indicates to a reset exception handler that a reset occurred if software had previously set this bit 1 Indicates to a reset exception handler that no reset occurred if software had previously set this bit Provided for software use - set anytime by software, cleared by reset.
16	—	Reserved ¹
17	TBEN ²	TimeBase Enable 0 TimeBase is disabled 1 TimeBase is enabled
18	Reserved ²	Reserved
19	DCLREE	Debug Interrupt Clears MSR[EE] 0 MSR[EE] unaffected by Debug Interrupt 1 MSR[EE] cleared by Debug Interrupt This bit controls whether Debug interrupts force External Input interrupts to be disabled, or whether they remain unaffected.
20	DCLRCE	Debug Interrupt Clears MSR[CE] 0 MSR[CE] unaffected by Debug Interrupt 1 MSR[CE] cleared by Debug Interrupt This bit controls whether Debug interrupts force Critical interrupts to be disabled, or whether they remain unaffected.
21	CICLRDE	Critical Interrupt Clears MSR[DE] 0 MSR[DE] unaffected by Critical class interrupt 1 MSR[DE] cleared by Critical class interrupt This bit controls whether certain Critical interrupts (Critical Input, Watchdog Timer) force Debug interrupts to be disabled, or whether they remain unaffected. Machine Check interrupts have a separate control bit. Note that if Critical Interrupt Debug events are enabled (DBCR0[CIRPT] set (which should only be done when the Debug APU is enabled), and MSR[DE] is set at the time of a Critical interrupt, a debug event is generated after the Critical Interrupt Handler has been fetched, and the Debug handler is executed first. In this case, DSRR0[DE] is cleared, such that after returning from the debug handler, the Critical interrupt handler is not run with MSR[DE] enabled.

Table 2-8. Hardware Implementation Dependent Register 0 (continued)

Bits	Name	Description
22	MCCLRDE	Machine Check Interrupt Clears MSR[DE] 0 MSR[DE] unaffected by Machine Check interrupt 1 MSR[DE] cleared by Machine Check interrupt This bit controls whether Machine Check interrupts force Debug interrupts to be disabled, or whether they remain unaffected. Note that if Critical Interrupt Debug events are enabled (DBCR0[CIRPT] set (which should only be done when the Debug APU is enabled), and MSR[DE] is set at the time of a Machine Check interrupt, a debug event is generated after the Machine Check interrupt handler has been fetched, and the Debug handler is executed first. In this case, DSRR0[DE] is cleared, such that after returning from the Debug handler, the Machine Check handler is not run with MSR[DE] enabled.
23	DAPUEN	Debug APU enable 0 Debug APU disabled 1 Debug APU enabled This bit controls whether the Debug APU is enabled. When enabled, Debug interrupts use the DSRR0/DSRR1 registers for saving state, and the se_rfdi instruction is available for returning from a debug interrupt. When disabled, Debug Interrupts use the critical interrupt resources CSRR0/CSRR1 for saving state, the se_rfdi instruction is used for returning from a debug interrupt, and the se_rfdi instruction is treated as an illegal instruction. When disabled, the settings of the DCLREE, DCLRCE, CICLRDE, and MCCLRDE bits are ignored and are assumed to be '1's Read and write access to DSRR0/DSRR1 via the mfspr and mtspr instructions is not affected by this bit.
24:31	—	Reserved ¹

¹ These bits are not implemented and should be written with zero for future compatibility.

² Unused on e200z0 and e200z0h

2.3.10 Hardware Implementation Dependent Register 1 (HID1)

The HID1 register is used for bus configuration and system control. HID1 is shown in [Figure 2-12](#).




Figure 2-12. Hardware Implementation Dependent Register 1 (HID1)

Table 2-9. Hardware Implementation Dependent Register 1

Bits	Name	Description
0:15	—	Reserved ¹
16:23	SYSCTL	System Control These bits are reflected on the outputs of the p_hid1_sysctl[0:7] output signals for use in controlling the system. They may need external synchronization.
24	ATS	Atomic status (read-only) Indicates state of the reservation bit in the load/store unit. See Section 3.4, “Memory Synchronization and Reservation Instructions,” for more detail.
25:31	—	Reserved ¹

¹ These bits are not implemented and should be written with zero for future compatibility.

The BUCSR register is used for general control and status of the branch target buffer (BTB). BUCSR is shown in [Figure 2-13](#).




Figure 2-13. Branch Unit Control and Status Register (BUCSR)

The BUCSR fields are defined in [Table 2-10](#).

Table 2-10. Branch Unit Control and Status Register

Bits	Name	Description
0:21 [32:53]	—	Reserved ¹
22 [54]	BBFI	Branch target buffer flash invalidate. When written to a ‘1’, BBFI flash clears the valid bit of all entries in the branch buffer; clearing occurs regardless of the value of the enable bit (BPEN). Note: BBFI is always read as 0.
25:30 [55:62]	—	Reserved ¹
31 [63]	BPEN	Branch target buffer enable. 0 Branch target buffer prediction disabled 1 Branch target buffer prediction enabled (enables BTB to predict branches) When the BPEN bit is cleared, no hits are generated from the BTB, and no new entries are allocated. Entries are not automatically invalidated when BPEN is cleared, the BBFI bit controls entry invalidation.

¹ These bits are not implemented and should be written with zero for future compatibility.

2.3.12 L1 Cache Configuration Register (L1CFG0)

The L1CFG0 register provides configuration information for an L1 cache supplied with this version of the e200 CPU core. For e200z0 and e200z0h reads of this register return a value of all zeros.

2.3.13 MMU Configuration Register (MMUCFG)

The MMUCFG register provides configuration information for the MMU supplied with this version of the e200 CPU core. For e200z0 and e200z0h, because no MMU is present, reads of this register return a value of 0x0000_0003, indicating MMU architecture version 3, and a null MMU.

2.4 SPR Register Access

SPRs are accessed with the **mf spr** and **mt spr** instructions. The following sections outline additional access requirements.

2.4.1 Invalid SPR References

System behavior when an invalid SPR is referenced depends on the apparent privilege level of the register. The register privilege level is determined by bit 5 in the SPR address. If the invalid SPR is accessible in user mode, then an illegal exception is generated. If the invalid SPR is accessible only in supervisor mode and the CPU core is in supervisor mode (MSR[PR] = 0), then an illegal exception is generated. If the invalid SPR address is accessible only in supervisor mode and the core is not in supervisor mode (MSR[PR] = 1), then a privilege exception is generated.

Table 2-11. System Response to Invalid SPR Reference

SPR Address Bit 5	Mode	MSR[PR]	Response
0	—	—	Illegal exception
1	supervisor	0	Illegal exception
1	user	1	Privilege exception

References to the SPRs associated with an optional unit (Cache, MMU, EFPU) when the unit is not present are treated as references to an invalid SPR unless otherwise defined.

2.4.2 Synchronization Requirements for SPRs

With the exception of the following registers, there are no synchronization requirements for accessing SPRs beyond those stated in Power Architecture Book E. Software requirements for synchronization before/after accessing these registers are shown in Table 2-12. The notation CSI in the table refers to a Context Synchronizing instruction which include **se_sc**, **isync**, **se_rfi**, **se_rfci**, and **se_rfdi**.

Table 2-12. Additional Synchronization Requirements for SPRs

Context Altering Event or Instruction		Required Before	Required After	Notes
mf spr				
DBSR	Debug Status Register	msync	none	
HID0	Hardware implementation dependent reg 0	none	none	
HID1	Hardware implementation dependent reg 1	msync	none	
mt spr				
BUCSR	Branch Unit Control and Status Register	none	CSI	
DBCR0	Debug Control Register 0	none	CSI	
DBCR1	Debug Control Register 1	none	CSI	
DBCR2	Debug Control Register 2	none	CSI	
DBSR	Debug Status Register	msync	none	
HID0	Hardware implementation dependent reg 0	CSI	CSI	
HID1	Hardware implementation dependent reg 1	none	CSI	
PID	PID0 register	none	CSI	

Note:

1. Not required if counter is not currently enabled

2.4.3 Special Purpose Register Summary

Power Architecture Book E and implementation-specific SPRs for the e200 core are listed in the following table. All registers are 32-bits in size. Register bits are numbered from bit 0 to bit 31 (most-significant to least-significant). An SPR register may be read or written with the **mf spr** and **mt spr** instructions. In the instruction syntax, compilers should recognize the mnemonic name given in the table below.

Table 2-13. Special Purpose Registers

Mnemonic	Name	SPR Number	Access	Privileged	e200-Specific
BUCSR	Branch Unit Control and Status Register	1013	R/W	Yes	Yes
CSRR0	Critical Save/Restore Register 0	58	R/W	Yes	No
CSRR1	Critical Save/Restore Register 1	59	R/W	Yes	No
CTR	Count Register	9	R/W	No	No

Table 2-13. Special Purpose Registers (continued)

Mnemonic	Name	SPR Number	Access	Privileged	e200-Specific
DAC1	Data Address Compare 1	316	R/W	Yes	No
DAC2	Data Address Compare 2	317	R/W	Yes	No
DBCR0	Debug Control Register 0	308	R/W	Yes	No
DBCR1	Debug Control Register 1	309	R/W	Yes	No
DBCR2	Debug Control Register 2	310	R/W	Yes	No
DBSR	Debug Status Register	304	Read/Clear ¹	Yes	No
DEAR	Data Exception Address Register	61	R/W	Yes	No
DSRR0	Debug save/restore register 0	574	R/W	Yes	Yes
DSRR1	Debug save/restore register 1	575	R/W	Yes	Yes
ESR	Exception Syndrome Register	62	R/W	Yes	No
HID0	Hardware implementation dependent reg 0	1008	R/W	Yes	Yes
HID1	Hardware implementation dependent reg 1	1009	R/W	Yes	Yes
IAC1	Instruction Address Compare 1	312	R/W	Yes	No
IAC2	Instruction Address Compare 2	313	R/W	Yes	No
IAC3	Instruction Address Compare 3	314	R/W	Yes	No
IAC4	Instruction Address Compare 4	315	R/W	Yes	No
IVPR	Interrupt Vector Prefix Register	63	R/W	Yes	No
LR	Link Register	8	R/W	No	No
L1CFG0	L1 cache config register 0	515	Read-only	No	Yes
MCSR	Machine Check Syndrome Register	572	R/W	Yes	Yes
MMUCFG	MMU configuration register	1015	Read-only	Yes	Yes
PID0	Process ID Register	48	R/W	Yes	No
PIR	Processor ID Register	286	Read-only	Yes	No
PVR	Processor Version Register	287	Read-only	Yes	No
SPRG0	SPR General 0	272	R/W	Yes	No
SPRG1	SPR General 1	273	R/W	Yes	No
SRR0	Save/Restore Register 0	26	R/W	Yes	No
SRR1	Save/Restore Register 1	27	R/W	Yes	No
SVR	System Version Register	1023	Read-only	Yes	Yes
XER	Integer Exception Register	1	R/W	No	No

Notes:

- ¹ The Debug Status Register can be read using **mfspr RT,DBSR**. The Debug Status Register cannot be directly written to. Instead, bits in the Debug Status Register corresponding to 1 bits in GPR(RS) can be cleared using **mtspr DBSR,RS**.

2.4.4 Reset Settings

Table 2-14 shows the state of the Power Architecture Book E architected registers and other optional resources immediately following a system reset.

Table 2-14. Reset Settings for e200 Resources

Resource	System Reset Setting
Program Counter	p_rstbase[0:29] 2'b00
GPRs	Unaffected ¹
CR	Unaffected ¹
BUCSR	0x0000_0000
CSRR0	Unaffected ¹
CSRR1	Unaffected ¹
CTR	Unaffected ¹
DAC1	0x0000_0000
DAC2	0x0000_0000
DBCR0	0x0000_0000
DBCR1	0x0000_0000
DBCR2	0x0000_0000
DBSR	0x1000_0000
DEAR	Unaffected ¹
DSRR0	Unaffected ¹
DSRR1	Unaffected ¹
ESR	0x0000_0000
HID0	0x0000_0000
HID1	0x0000_0000
IAC1	0x0000_0000
IAC2	0x0000_0000
IAC3	0x0000_0000
IAC4	0x0000_0000
IVPR	Unaffected ¹
LR	Unaffected ¹
L1CFG0 ²	—
MCSR	0x0000_0000
MMUCFG ²	—
MSR	0x0000_0000
PID0	0x0000_0000

Table 2-14. Reset Settings for e200 Resources (continued)

Resource	System Reset Setting
PIR ²	—
PVR ²	—
SPRG0	Unaffected ¹
SPRG1	Unaffected ¹
SRR0	Unaffected ¹
SRR1	Unaffected ¹
SVR ²	—
XER	0x0000_0000

¹ Undefined on *m_por* assertion, unchanged on *p_reset_b* assertion

² Read-only register

Chapter 3

Instruction Model

This chapter provides additional information about the Power Architecture Book E as it relates specifically to e200.

The e200z0 and e200z0h cores are a 32-bit implementation of the Power Architecture Book E as defined in Power Architecture Book E Specification v 2.0. This architecture specification includes a recognition that different processor implementations may require clarifications, extensions or deviations from the architectural descriptions. e200z0 and e200z0h are unique in that they support only the VLE instruction set encodings. The VLE APU is described in *PowerPC VLE, version 1.01*.

3.1 Unsupported Instructions and Instruction Forms

The e200 core does not support the instructions listed in [Table 3-1](#). An unimplemented instruction exception is generated if the processor attempts to execute one of these instructions.

Table 3-1. List of Unsupported Instructions

Type/Name	Mnemonics
String Instructions	lswi, lswx, stswi, stswx
Device control register and Move from APID	mfapidi, mfdcrx, mtdcrx

3.2 Optionally Supported Instructions and Instruction Forms

e200 cores optionally supports the instructions listed in [Table 3-2](#) if a cache and/or TLB is present. An instruction exception may be generated if the processor attempts to execute one of these instructions and the related functional block is not present, or the specific instruction may be treated as a no-op.

Table 3-2. List of Optionally Supported Instructions

Type/Name	Mnemonics	Unit
Cache Management Instructions ¹	dcba, dcbf, dcbi, dcbt, dcbtst, dcbst, dcbz, icbi, icbt	Data Cache/ Unified Cache Instruction Cache/Unified Cache
Cache Locking Instructions ²	dcbtls, dcibtstls, dcbic, icbtls, icblc	Data Cache/ Unified Cache Instruction Cache/Unified Cache
TLB Management Instructions ³	tlbivax, tlbre, tlbsx, tlbsync, tlwwe	TLB
DCR Management ³	mfdcr, mtdcr	DCR

¹ These instructions are not supported and are treated as no-ops, with the exception of **dcbz** which results in an Alignment Interrupt, and **dcbi**, which is treated as a privileged no-op.

² These instructions are not supported and are treated as no-ops.

³ These instructions are not supported and are treated as unimplemented.

3.3 Memory Access Alignment Support

The e200 core provides hardware support for unaligned memory accesses; however, there is a performance degradation for accesses that cross a 32-bit (4-byte) boundary. For these cases, the throughput of the load/store unit is degraded to 1 misaligned load every 2 cycles. Stores that are misaligned across a 32-bit (4-byte) boundary can be translated at a rate of 2 cycles per store. Frequent use of unaligned memory accesses result in an impact on performance.

NOTE

Accesses that cross a 32-bit boundary may be restarted.

3.4 Memory Synchronization and Reservation Instructions

The **msync** instruction provides a synchronization function and a memory barrier function. This instruction waits for all preceding instructions and data memory accesses to complete before the **msync** instruction completes. Subsequent instructions in the instruction stream are not initiated until after the **msync** instruction ensures these functions have been performed.

On the e200 core, the **mbar** instruction behaves identically to the **msync** instruction. The **mbar** instruction MO field is ignored by the e200 core.

The e200 core implements the **lwarx** and **stwcx** instructions as described in Book E. If the EA is not a multiple of 4 for either instruction, an alignment interrupt is invoked.

As allowed by Power Architecture Book E, the e200 core does not require that for a **stwcx** instruction to succeed, the EA of the **stwcx** instruction must be to the same reservation granule as the EA of a preceding **lwarx** instruction. Reservation granularity is implementation-dependent. The e200 core does not define a reservation granule explicitly; reservation granularity is defined by external logic. When no external logic is provided, the e200 core performs no address comparison checking, thus the effective implementation granularity is “null”.

The e200 core implements an internal reservation status flag (HID1[ATS]) representing reservation status. This flag is set when a **lwarx** instruction is executed and completes without error, and remains set until it is cleared by one of the following mechanisms:

- Execution of a **stwcx** instruction is completed without error, or
- The e200 core **p_rsrv_clr** input signal is asserted, or
- The reservation is invalidated when an external input, critical input (on e200z0), or non-maskable interrupt (on e200z0h) is signaled and the HID0[ICR] bit is set.

When the e200 core decodes a **stwcx** instruction, it checks the value of the local reservation flag (HID1[ATS]). If the status indicates that no reservation is active, then the **stwcx** instruction is treated as a nop. No exceptions are taken, and no access is performed, thus no data breakpoint occurs, regardless of matching the data breakpoint attributes.

The e200 core provides the input signal *p_hresp*[2:0], which is sampled at termination of a **stwcx** store transfer to allow an external agent or mechanism to indicate that the **stwcx** instruction has failed to update memory, even though a reservation existed for the store at the time it was issued. This is not considered an

error, and causes the condition codes for the **stwcx.** instruction to be written as if a reservation did not exist for the **stwcx.** instruction. In addition, any outstanding reservation is cleared.

The **p_rsrv_clr** input signal is not intended for normal use in managing reservations. It is provided for specialized system applications. The normal bus protocol is used to manage reservations using external reservation logic in systems with multiple coherent bus masters, using the transfer type and transfer response signals. In single coherent master systems, no external logic is required, and the internal reservation flag is sufficient to support multi-tasking applications.

3.5 Branch Prediction

In the e200z0h, the instruction fetching mechanism uses a branch target buffer to detect branch instructions early. This branch instruction lookahead scheme allows branch targets to be fetched early, thereby hiding some taken branch bubbles.

3.6 Interruption of Instructions by Interrupt Requests

In general, the e200 core samples pending non-maskable interrupts, external input, and critical input interrupt requests at instruction boundaries. However, in order to reduce interrupt latency, long running instructions may be interrupted prior to completion. Instructions in this class include divides (**divw[uo][.I]**), load multiple word and store multiple word. When interrupted prior to completion, the value saved in SRR0/CSRR0 is the address of the interrupted instruction. The instruction is restarted from the beginning after returning from the interrupt handler.

3.7 New e200 Instructions

The e200 core implements the Freescale EIS **isel** APU as described below which extends the Power Architecture Book E instruction set. The e200 **wait** instruction implements a wait for interrupt function and is described below. The e200 **se_rfdi** instruction returns from a Debug interrupt and is also described below.


3.7.1 ISEL APU

The ISEL APU defines the **isel** instruction which provides a means to select one of two registers and place the result in a destination register under the control of a predicate value supplied by a bit in the condition register. This instruction can be used to eliminate branches in software and in many cases improve performance. This instruction can also increase program execution time determinism by eliminating the need to predict the target and direction of the branches replaced by the integer select function. The instruction form and definition is as follows.

isel

Integer Select

isel RT, RA, RB, crb



```

if RA=0 then a ← 320 else a ← GPR(RA)
c = CRcrb
if c then GPR(RT) ← a
else GPR(RT) ← GPR(RB)

```

For **isel**, if the bit of the CR specified by (crb) is set, the contents of RA|0 are copied into RT. If the bit of the CR specified by (crb) is clear, the contents of RB are copied into RT.

Other registers altered:

- None

3.7.2 Debug APU

e200 implements the EIS Debug APU, as documented in the *EREF: A Programmer's Reference Manual for Freescale Book E Processors*, to support the capability to handle the Debug interrupt as an additional interrupt level. To support this interrupt level, a new “return from debug interrupt” (**se_rfdi**) instruction is defined as part of the Debug APU, along with a new pair of save/restore registers, DSRR0, and DSRR1.


When the Debug APU is enabled (HID0[DAPUEN] = 1), the **se_rfdi** instruction provides a means to return from a debug interrupt. See [Section 2.3.9, “Hardware Implementation Dependent Register 0 \(HID0\),”](#) for more information about enabling the Debug APU.

The instruction forms and definition are as follows.

se_rfdi

Return From Debug Interrupt

se_rfdi



```

MSR ← DSRR1
PC ← DSRR032:62 || 0b0

```

se_rfdi

The **se_rfdi** instruction is used to return from a Debug interrupt, or as a means of simultaneously establishing a new context and synchronizing on that new context.

The contents of Debug Save/Restore Register 1 are place into the Machine State Register. If the new Machine State Register value does not enable any pending exceptions, then the next instruction is fetched, under control of the new Machine State Register value from the address DSRR0_{32:62}|| 0b0. If the new Machine State Register value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into Save/Restore Register 0 or Critical Save/Restore Register 0 by the interrupt processing mechanism is the address of the instruction that would have been executed next had the interrupt not occurred (such as the address in Debug Save/Restore Register 0 at the time of the execution of the **se_rfdi**).

Execution of this instruction is privileged and context synchronizing.

Special Registers Altered:

- MSR

When the Debug APU is disabled (HID0[DAPUEN]=0), this instruction is treated as an illegal instruction.

3.7.3 WAIT APU


The **wait** instruction allows software to cease all synchronous activity, waiting for an asynchronous interrupt to occur. The instruction can be used to cease processor activity in both user and supervisor modes. Asynchronous interrupts which cause the waiting state to be exited if enabled are critical input, external input, machine check pin (*p_mcp_b*). Nonmaskable interrupts (*p_nmi_b*) also cause the waiting state to be exited.

wait

Wait for Interrupt

wait

wait



The **wait** instruction provides an ordering function for the effects of all instructions executed by the processor executing the **wait** instruction and stops synchronous processor activity. Executing a **wait** instruction ensures that all instructions have completed before the **wait** instruction completes, causes processor instruction fetching to cease, and ensures that no subsequent instructions are initiated until an asynchronous interrupt or a debug interrupt occurs.

Once the **wait** instruction has completed, the program counter points to the next sequential instruction. The saved value in xSRR0 when the processor re-initiates activity points to the instruction following the **wait** instruction.

Execution of a wait instruction places the CPU in the “waiting” state and is indicated by assertion of the **p_waiting** output signal. The signal is negated after leaving the “waiting” state.

Software must ensure that interrupts responsible for exiting the waiting state are enabled before executing a wait instruction.

3.8 Unimplemented SPRs and Read-Only SPRs

e200 fully decodes the SPR field of the **mfspir** and **mtspir** instructions. If the SPR specified is undefined and not privileged, an illegal instruction exception is generated. If the SPR specified is undefined and privileged and the CPU is in user mode (MSR[PR=1]), a privileged instruction exception is generated. If the SPR specified is undefined and privileged and the core is in supervisor mode (MSR[PR=0]), an illegal instruction exception is generated.

For the **mtspir** instruction, if the SPR specified is read-only and not privileged, an illegal instruction exception is generated. If the SPR specified is read-only and privileged and the core is in user mode (MSR[PR=1]), a privileged instruction exception is generated. If the SPR specified is read-only and privileged and the core is in supervisor mode (MSR[PR=0]), an illegal instruction exception is generated.

For e200z0 and e200z0h, the following SPRs are not implemented and attempted access via a **mtspir** or **mfspir** instruction results in an unimplemented instruction exception, unless the register is privileged and the access attempt is made in user mode, in which case a privileged instruction exception occurs.

Table 3-3. List of Unimplemented SPRs

Type	Name
Timebase	DEC, DECAR, TCR, TSR, TBU, TBL
Software-Use Special Purpose Registers	USPRG0, SPRG2-7
Interrupt Vector Offset Registers	IVOR0-15 ¹

¹ These SPRs are hardwired to specific values, and are readable, but a **mtspir** results in an unimplemented or privileged exception.

3.9 Invalid Forms of Instructions

3.9.1 Load and Store with Update Instructions

Power Architecture Book E defines the case when a load with update instruction specifies the same register in the RT and RA field of the instruction as an invalid format. For this invalid case, the e200 core performs the instruction and update the register with the load data. In addition, if RA=0 for any load or store with update instruction, the e200 core updates RA (GPR0).

3.9.2 Load Multiple Word (e_lmw) Instruction

Power Architecture Book E defines as invalid any form of the **e_lmw** instruction in which RA is in the range of registers to be loaded, including the case in which RA=0. On e200, invalid forms of the **e_lmw** instruction is executed as follows:

- Case 1: *RA is in the range of RT, RA!=0*. In this case, address generation for individual loads to register targets in is done using the architectural value of RA which existed when beginning

execution of this **e_lmw** instruction. RA is overwritten with a value fetched from memory as if it had not been the base register. Note that if the instruction is interrupted and restarted, the base address may be different if RA has been overwritten.

- Case 2: *RA=0 and RT=0*. In this case, address generation for all loads to register targets RT=0 to RT=31 is done substituting the value of 0 for the RA operand.

3.9.3 Instructions with Reserved Fields Non-Zero

Power Architecture Book E defines certain bit fields in various instructions as reserved and specifies that these fields be set to zero. Per the Book E recommendation, e200 ignores the value of the reserved field (bit 31) in X-form integer load and store instructions. For all other instructions, e200 generates an illegal instruction exception if a reserved field is non-zero.

3.10 Optionally Supported APU Instructions

e200 cores optionally support several APUs. If a core does not implement a particular APU, it may treat these instructions as illegal, or may treat them as unimplemented. e200z0 and e200z0h treat the Embedded Floating-Point APU instructions (efs_{xxx}, brinc) as unimplemented instructions because other cores of the e200 family implement this APU. All other non-supported APUs are treated as illegal instructions.



Chapter 4

Instruction Pipeline and Execution Timing

This section describes the e200 instruction pipeline and instruction timing information. The core is partitioned into the following subsystems:

- Instruction unit
- Control unit
- Integer unit
- Load/store unit
- Core interface

4.1 Overview of Operation

A block diagram of the e200 cores are shown in [Figure 4-2](#) and [Figure 4-1](#). The instruction fetch unit prefetches instructions from memory into the instruction buffers. The decode unit decodes each instruction and generates information needed by the branch unit and the execution units.

The instruction fetch unit attempts to supply a constant stream of instructions to the execution pipeline. In the e200z0h it does so by decoding and detecting branches early in the instruction buffer, making branch predictions, and prefetching their branch targets into the instruction buffer. By prefetching the branch targets early, some or all of the branch pipeline bubbles can be hidden from the execution pipeline.

The instruction issue unit attempts to issue a single instruction each cycle to one of the execution units. Source operands for each of the instructions are provided from the GPRs or from the operand feed-forward muxes. Data or resource hazards may create stall conditions which cause instruction issue to be stalled for one or more cycles until the hazard is eliminated.

The execution units write the result of a finished instruction onto the proper result bus and into the destination registers. The writeback logic retires an instruction when the instruction has finished execution. Up to two results can be simultaneously written.




Figure 4-1. e200z0h Block Diagram




Figure 4-2. e200z0 Block Diagram

4.1.1 Control Unit

The control unit coordinates the instruction fetch unit, branch unit, instruction decode unit, instruction issue unit, completion unit and exception handling logic.

4.1.2 Instruction Unit

The instruction unit controls the flow of instructions to the instruction buffers and decode unit. A set of instruction prefetch buffers allow the instruction unit to fetch instructions ahead of actual execution, and serve to decouple memory and the execution pipeline.

4.1.3 Branch Unit

In the e200z0h the branch unit contains a single entry Branch Target Buffer (BTB) to accelerate execution of branch instructions.

Conditional branches which are not taken execute in a single clock on e200. Branches with successful target prefetching have an effective execution time of one clock in the e200z0h. All other taken branches have an execution time of two clocks on e200.

4.1.4 Instruction Decode Unit

The decode unit includes the instruction buffers. A single instruction can be decoded each cycle. The major functions of the decode logic are:

- Opcode decoding to determine the instruction class and resource requirements for each instruction being decoded.
- Source and destination register dependency checking.
- Execution unit assignment.
- Determine any decode serializations, and inhibit subsequent instruction decoding.

The decode unit operates in a single processor clock cycle.

4.1.5 Exception Handling

The exception handling unit includes logic to handle exceptions, interrupts, and traps.

4.2 Execution Units

The core data execution units consist of the integer unit, and the load/store unit. Included in the execution units section are the 32 general purpose registers (GPRs). Instructions with data dependencies begin execution when all such dependencies are resolved.

4.2.1 Integer Execution Unit

The integer execution unit is used to process arithmetic and logical instructions. Adds, subtracts, compares, count leading zeros, shifts and rotates execute in a single cycle.

Multiply instructions have a data-dependent latency and throughput rate of 1–4 cycles.

Divide instructions have a latency of 5–34 cycles depending on the operand data. While the divide is running, the rest of the pipeline is unavailable for additional instructions (blocking divide).

4.2.2 Load/Store Unit

The load/store unit executes instructions that move data between the GPRs and the memory subsystem. A load followed by a dependent instruction does not incur any pipeline stall, except when the dependent instruction is a load/store instruction, and the latter instruction is using the previous load data for its effective address (EA) calculation, in which case a 1 cycle “register-busy” pipeline stall is incurred.

Loads, when free of the above effective address calculation dependency, execute with a maximum throughput of one per cycle and one cycle latency. Store data can be fed-forward from an immediately preceding load with no stall.

4.3 Instruction Pipeline

The four stage processor pipeline consists of stages for instruction fetch (IFETCH), instruction decode (DECODE), execution (EXECUTE), and result writeback (WB). For memory operations, the effective address generation occurs in the decode stage, while the memory access occurs in the execute stage.

The processor also contains an instruction prefetch buffer to allow buffering of instructions prior to the decode stage. Instructions proceed from this buffer to the instruction decode stage by entering the instruction decode register IR.

Table 4-1. Pipeline Stages

Stage	Description
IFETCH	Instruction Fetch From Memory
DECODE/EA	Instruction Decode / Register Read/ Operand Forwarding / EA Calculation
EXECUTE/MEM	Instruction Execution / Memory Access
WB	Write Back to Registers




Figure 4-3. Pipeline Diagram

4.3.1 Description of Pipeline Stages

The IFetch pipeline stage retrieve instructions from the memory system and determine where the next instruction fetch is performed. Up to two 16-bit instructions are sent from memory to the instruction buffers every cycle.

The Decode pipeline stage decodes instructions and performs dependency checking. Simple integer instructions complete execution in the Execute stage of the pipeline.

Execution of load/store instructions is pipelined. The effective address calculations for load/store instructions are performed in the decode stage. This effective address is driven out to the data memory in the same stage. The actual memory access occurs in the execute stage.

Load-to-use dependencies do not incur pipeline bubbles except when the dependent instruction is a load or store instruction, and the latter instruction is dependent on its previous load data for EA calculation. If an ALU instruction is dependent on a load instruction, the data is fed directly into the ALU for execution. No pipeline bubble is incurred in this case.

Multiply instructions require 1 to 4 clocks to execute.

All condition-setting instructions complete in the Execute stage of the pipeline.

Result feed-forward hardware forwards the result of one instruction into the source operand(s) of a following instruction so that the execution of data-dependent instructions do not wait until the completion of the result writeback.

4.3.2 Instruction Buffers

e200 contains a set of instruction buffers which supply instructions into the Instruction Register (IR) for decoding.

In normal sequential execution, instructions are loaded into the IR from Slot 0, and whenever a slot is empty, a 32-bit prefetch is initiated which fills the earliest empty slot beginning with Slot 0.

If the instruction buffer empties, instruction issue stalls, and the buffer is refilled. The first returned instruction is forwarded directly to the IR.




Figure 4-4. e200 Instruction Buffers




Figure 4-5. e200 Instruction Buffers

4.3.2.1 Branch Prediction in e200z0h

In e200z0h the HID0[BPRED] field is used to control whether prediction will be made for forward or backward branches (or both).

To resolve branch instructions and improve the accuracy of branch predictions, e200 implements a dynamic branch prediction mechanism using a 1-entry branch target buffer (BTB), a fully associative address cache of branch target addresses. The BTB on e200 is purposefully small to reduce cost and power. It is expected to accelerate the execution of loops.

An entry is allocated in the BTB whenever a branch resolves as taken and the BTB is enabled. Branches that have not been allocated are always predicted as not taken. Entries in the BTB are allocated on taken branches using a FIFO replacement algorithm.

Each BTB entry holds a 2-bit branch history counter, whose value is incremented or decremented on a BTB hit, depending on whether the branch was taken. The counter can assume four different values: strongly taken, weakly taken, weakly not taken, and strongly not taken.

A branch will be predicted as taken on a hit in the BTB with a counter value of strongly or weakly taken. In this case the target address contained in the BTB is used to redirect the instruction fetch stream to the target of the branch prior to the branch reaching the instruction decode stage. In the case of a mispredicted branch, the instruction fetch stream returns to the sequential instruction stream after the branch has been resolved.

When a branch is predicted taken and the branch is later resolved (in the branch decode stage), the value of the counter is updated. A branch whose counter indicates weakly taken is resolved as taken, the counter increments so that the prediction becomes strongly taken. If the branch resolves as not taken, the prediction changes to weakly not-taken. The counter saturates in the strongly taken states when the prediction is correct.

e200 does not implement the static branch prediction that is defined by the PowerPC architecture. The BO prediction bit in branch encodings is ignored.

Dynamic branch prediction is enabled by setting BUCSR[BPEN]. Clearing BUCSR[BPEN] disables dynamic branch prediction, in which case e200 predicts every branch as not taken. Additional control is available in the HID0[BPRED] field to control whether forward or backward branches (or both) are candidates for entry into the BTB, and thus for branch prediction. Once a branch is in the BTB, HID0[BPRED] has no further effect on that branch entry.

The BTB uses virtual addresses for performing tag comparisons. On allocation of a BTB entry, the effective address of a taken branch, along with the current Instruction Space (as indicated by MSR[IS]) is loaded into the entry and the counter value is set to weakly taken. The current PID value is not maintained as part of the tag information.

e200 does support automatic flushing of the BTB when the current PID value is updated by a **mtcr PID0** instruction. Software is otherwise responsible for maintaining coherency in the BTB when a change in effective to real (virtual to physical) address mapping is changed. This is supported by the BUCSR[BBFI] control bit.



IS = Instruction Space

Figure 4-6. e200 Branch Target Buffer

4.3.3 Single-Cycle Instruction Pipeline Operation

Sequences of single-cycle execution instructions follow the flow in Figure 4-7. Instructions are issued and completed in program order. Most arithmetic and logical instructions fall into this category.




Figure 4-7. Basic Pipeline Flow, Single Cycle Instructions

4.3.4 Basic Load and Store Instruction Pipeline Operation

The effective address (EA) calculations for load and store instructions are performed in the decode stage. The memory access occurs in the execution stage.

If a load instruction is followed by an dependent ALU instruction, the load data is driven from the memory in the MEM stage and feed-forwarded into the dependent ALU instruction in the following cycle. As a result, there is no load-to-use pipeline bubble. Figure 4-8 shows the instruction flow for a load instruction followed by a dependent add instruction.




Figure 4-8. A Load Followed By A Dependent Add Instruction

Back-to-back load/store instructions are executed in a pipelined fashion, provided that their effective address calculations are not dependent on their previous load instructions. Figure 4-9 shows the basic pipe line flow for two back-to-back load instructions. In this case, the 2nd load does not depend on its previous load data for its EA calculation. Notice that the memory access of the first load instruction overlaps in time with the EA calculation of the second load instruction.




Figure 4-9. Back-to-back Load Instructions

When a load is followed by a load or a store instruction that depends on the first load data for EA calculation, a pipeline stall is incurred. Figure 4-10 shows the instruction flow for a load instruction followed by a dependent store instruction through EA calculation. The second store instruction, in this case, is dependent on the first load instruction for its EA calculation.




Figure 4-10. A Load Followed By A Dependent Store Instruction

A store instruction that depends on its previous load for its store data does not stall the pipeline.

4.3.5 Change-of-Flow Instruction Pipeline Operation

A branch instruction takes either one or two cycles to execute. Simple change of flow instructions require 2 cycles to refill the pipeline with the target instruction for taken branches and branch and link instructions with no prediction.




Figure 4-11. Basic Pipeline Flow, Branch Instructions

For branch type instructions in e200z0h, in some situations this 2 cycle timing may be reduced by performing the target fetch speculatively while the branch instruction is still being fetched into the instruction buffer. The branch target address is obtained from the BTB. The resulting branch timing reduces to a single clock when the target fetch is initiated early enough and the branch is taken.




Figure 4-12. Basic Pipeline Flow, Branch Speculation

4.3.6 Basic Multi-Cycle Instruction Pipeline Operation

The multiply, divide, and load and store multiple instructions require multiple cycles in the execute stage.




Figure 4-13. Basic Pipeline Flow, Multi-cycle Instructions

Instructions must complete and write back results in order. A single cycle instruction which follows a multi-cycle instruction must wait for completion of the multi-cycle instruction prior to its writeback in order to meet the in-order requirement. Result feed-forward paths are provided so that execution may continue prior to result writeback.

4.3.7 Additional Examples of Instruction Pipeline Operation for Load and Store

Figure 4-14 shows an example of pipelining two non-data dependent load or store instructions with a following data dependent single cycle instruction. While the first load or store begins accessing memory in the MEM stage, the next load or store can be calculating a new effective address in the DEC/EA stage. The **add** in this example does not stall even though there is a data dependency on its preceding load instruction.




Figure 4-14. Pipelined Load/Store Instructions

For memory access instructions, wait-states may occur. This causes a following memory access instruction to stall because the following memory access may not be initiated as shown in [Figure 4-15](#). Here, the first **ld/st** instruction incurs a wait-state on the bus interface, causing succeeding instructions to stall.




Figure 4-15. Pipelined Load/Store Instructions with Wait-state

4.3.8 Move to/from SPR Instruction Pipeline Operation

Most **mtspr** and **mfspir** instructions are treated like single cycle instructions in the pipeline, and do not cause stalls. Exceptions are for the MSR, and the Debug SPRs that do cause stalls. The following figures show examples of **mtspr** and **mfspir** instruction timing.

[Figure 4-16](#) applies to the Debug SPRs. These instructions do not begin execution until all previous instructions have finished their execute stage. If the previous instruction of **mfspir** or **mtspr** is a multicycle instruction, the **mfspir** and **mtspr** instructions do not begin execution until its previous instruction moves into the WB stage as shown in [Figure 4-16](#). In addition, execution of subsequent instructions is stalled until the **mfspir** and **mtspr** instructions complete.




Figure 4-16. mtspr, mfspir Instruction Execution—(1)

Figure 4-17 applies to the **mtmsr** instruction and the **wrttee** and **wrtteei** instructions. Execution of subsequent instructions is stalled until these instructions writeback.




Figure 4-17. **mtmsr**, **wrttee**, **wrtteei** Instruction Execution

4.4 Control Hazards

Several internal control hazards exist in e200 which can cause certain instruction sequences to incur one or more stall cycles. These include the following:

- **mfspir** instruction preceded by a **mtspir** instruction—issue stalls until the **mtspir** completes

4.5 Instruction Serialization

The types of serialization required by the core are as follows:

- Completion serialization
- Dispatch (Decode/Issue) serialization
- Refetch serialization

4.5.1 Completion Serialization

A completion serialized instruction is held for execution until all prior instructions have completed. The instruction then executes after it is next to complete in program order. Results from these instructions are not available for or forwarded to subsequent instructions until the instruction completes. Instructions which are completion serialized are:

- Instructions that access or modify system control or status registers. for example, **mcrxr**, **mtmsr**, **wrttee**, **wrtteei**, **mtspir**, **mfspir** (except to CTR/LR),
- Instructions defined by the architecture as context or execution synchronizing: **se_isync**, **msync**, **se_rfi**, **se_rfdi**, **se_sc**.

4.5.2 Dispatch Serialization

Some instructions are dispatch-serialized by the core. An instruction that is dispatch-serialized prevents the next instruction from decoding until all instructions up to and including the dispatch-serialized instruction completes. Instructions which are dispatch serialized are **se_isync**, **mbar**, **msync**, **se_rfi**, **se_rfci**, **se_rfdi**, **se_sc**.

4.5.3 Refetch Serialization

Refetch serialized instructions inhibit dispatching of subsequent instructions and force a pipeline refill to refetch subsequent instructions after completion. These include:

- The context synchronizing instruction **isync**.
- The **se_rfi**, **se_rfci**, **se_rfdi**, and **se_sc** instructions.

4.6 Interrupt Recognition and Exception Processing

[Figure 4-18](#) shows timing for interrupt recognition and exception processing overhead. This example shows best-case response timing when an interrupt is received and processed during execution of a sequence of single-cycle instructions.





Figure 4-18. Interrupt Recognition and Handler Instruction Execution

Figure 4-19 shows timing for interrupt recognition and exception processing overhead. This example shows best-case response timing when an interrupt is received and processed during execution of a load or store instruction. The fetch for the handler is delayed until completion of the load or store, regardless of the number of wait-states.



**Figure 4-19. Interrupt Recognition and Handler Instruction Execution—
Load/Store in Progress**

Figure 4-20 shows timing for interrupt recognition and exception processing overhead. This example shows best-case response timing when an interrupt is received and processed during execution of a multicycle interruptible instruction.




Figure 4-20. Interrupt Recognition and Handler Instruction Execution—
Multi-Cycle Instruction Abort

4.7 Instruction Timings

Instruction timing in number of processor clock cycles for various instruction classes is shown in [Table 4-2](#). Pipelined instructions are shown with cycles of total latency and throughput cycles. Divide and multiply instructions are not pipelined and block other instructions from executing during execution.

Load/store multiple instruction cycles are represented as a fixed number of cycles plus a variable number of cycles where ‘n’ is the number of words accessed by the instruction. In addition, cycle times marked with an ‘&’ require variable number of additional cycles due to serialization.

Table 4-2. Instruction Class Cycle Counts

Class of Instructions	Latency	Throughput	Special Notes
integer: add, sub, shift, rotate, logical, cntlzw class instructions	1	1	—
integer: compare	1	1	—
Branch	2/1	2/1	Branches take either 2 or 1 cycles to execute. (Taken/Not Taken)
multiply	1-4	1-4	data dependent timing
divide	5-34	5-34	data dependent timing
CR logical	1	1	—
loads (non-multiple)	1	1	—
load multiple	1 + n	1 + n	Actual timing depends on n and address alignment. (n = number of registers transferred)
stores (non-multiple)	1	1	—
store multiple	1 + n	1 + n	Actual timing depends on n and address alignment. (n = number of registers transferred)
mtmsr, wrtee, wrteei	2&	2	—
mcrf	1	1	—
mf spr, mt spr	2&	2&	applies to Debug SPRs, optional unit SPRS
mf spr, mfmsr	1	1	applies to internal, non Debug SPRs
mfc r, mtc r	1	1	—
se_rfi, se_rfci, se_rfdi	3	—	—
se_sc	3	—	—
tw	3	—	Trap taken timing

4.8 Operand Placement on Performance

The placement (location and alignment) of operands in memory affects relative performance of memory accesses, and in some cases, affects it significantly. [Table 4-3](#) indicates the effects for the e200 core.

In [Table 4-3](#), optimal means that one effective address (EA) calculation occurs during the memory operation. Good means that multiple EA calculations occur during the memory operation which may cause additional bus activities with multiple bus transfers. Poor means that an alignment interrupt is generated by the storage operation.

Table 4-3. Performance Effects of Storage Operand Placement

Operand		
Size	Byte Align	Performance
4 Byte	4 <4	optimal good
2 Byte	2 <2	optimal good
1 Byte	1	optimal
lmw, stmw	4 <4	good poor
string	N/A	—

Notes:

Optimal: One EA calculation occurs.

Good: Multiple EA calculations occur which may cause additional bus activities with multiple bus transfers.

Poor: Alignment Interrupt occurs.

Chapter 5

Interrupts and Exceptions

The Power Architecture Book E document defines the mechanisms by which the e200 core implements interrupts and exceptions. The document uses the terminology ‘interrupt’ as the action in which the processor saves its old context and begins execution at a pre-determined interrupt handler address. ‘Exceptions’ are referred to as events which, when enabled, cause the processor to take an interrupt. This section uses the same terminology.

The Power Architecture exception mechanism allows the processor to change to supervisor state as a result of unusual conditions arising in the execution of instructions, and from external signals, bus errors, or various internal conditions. When interrupts occur, information about the state of the processor is saved to machine state save/restore registers (SRR0/SRR1, CSRR0/CSRR1, or DSRR0/DSRR1) and the processor begins execution at an address (interrupt vector) determined by the Interrupt Vector Prefix register (IVPR), and one of the hardwired Interrupt Vector Offset values. Processing of instructions within the interrupt handler begins in supervisor mode.

Multiple exception conditions can map to a single interrupt vector, and may be distinguished by examining registers associated with the interrupt. The Exception Syndrome register (ESR) is updated with information specific to the exception type when an interrupt occurs.

To prevent loss of state information, interrupt handlers must save the information stored in the machine state save/restore registers, soon after the interrupt has been taken. Three sets of these registers are implemented; SRR0 and SRR1 for non-critical interrupts, CSRR0 and CSRR1 for critical interrupts, and DSRR0 and DSRR1 for debug interrupts (when the Debug APU is enabled). Hardware supports nesting of critical interrupts within non-critical interrupts, and debug interrupts within both critical and non-critical interrupts. It is up to the interrupt handler to save necessary state information if interrupts of a given class are re-enabled within the handler.

The following terms are used to describe the stages of exception processing:

Recognition	Exception recognition occurs when the condition that can cause an exception is identified by the processor. This is also referred to as an exception event.
Taken	An interrupt is said to be taken when control of instruction execution is passed to the interrupt handler; that is, the context is saved and the instruction at the appropriate vector offset is fetched and the interrupt handler routine begins.
Handling	Interrupt handling is performed by the software linked to the appropriate vector offset. Interrupt handling is begun in supervisor mode.

Returning from an interrupt is performed by executing an **se_rfi**, **se_rfci**, or **se_rfdi** instruction to restore state information from the respective machine state save/restore register pair.

5.1 e200 Interrupts

As specified by the Power Architecture Book E specification, interrupts can be either precise or imprecise, synchronous or asynchronous, and critical or non-critical. Asynchronous exceptions are caused by events external to the processor's instruction execution; synchronous exceptions are directly caused by instructions or an event somehow synchronous to the program flow, such as a context switch. A precise interrupt architecturally guarantees that no instruction beyond the instruction causing the exception has (visibly) executed. Critical interrupts are provided with a separate save/restore register pair (CSRR0/CSRR1) to allow certain critical exceptions to be handled within a non-critical interrupt handler.

The types of interrupts handled are shown in [Table 5-1](#).

Table 5-1. Interrupt Classifications

Interrupt Types	Synchronous/Asynchronous	Precise/Imprecise	Critical/Non-Critical/Debug
System Reset	Asynchronous, non-maskable	Imprecise	—
Machine Check	—	—	Critical
Critical Input Interrupt	Asynchronous, maskable	Imprecise	Critical
External Input Interrupt	Asynchronous, maskable	Imprecise	Non-critical
Instruction-based Debug Interrupts	Synchronous	Precise	Critical / Debug
Debug Interrupt (UDE) Debug Imprecise Interrupt	Asynchronous	Imprecise	Critical / Debug
Data Storage / Alignment Interrupts Instruction Storage Interrupts	Synchronous	Precise	Non-critical

These classifications are discussed in greater detail in [Section 5.7, “Interrupt Definitions.”](#) Interrupts implemented in e200 and the exception conditions that cause them are listed in [Table 5-2](#).

Table 5-2. Exceptions and Conditions

Interrupt Type	Corresponding Interrupt Vector Offset	Causing Conditions
System reset	none, vector to $[p_rstbase[0:29]] \parallel 2'b00$	Reset by assertion of p_reset_b Debug Reset Control
Critical Input	IVOR 0 ¹	$p_critint_b$ is asserted and $MSR[CE]=1$
Machine check	IVOR 1	p_mcp_b is asserted and $MSR[ME] = 1$ 1. Bus error (XTE) with $MSR[EE]=0$ and current $MSR[ME]=1$ Non-maskable interrupt (p_nmi_b recognized asserted) regardless of $MSR[ME]$
Data Storage	IVOR 2	Access control. (unused on Zen Z0n2p and Zen Z0Hn2p) Precise external termination error (p_tea_b assertion and precise recognition) and $MSR[EE]=1$
Instruction Storage	IVOR 3	Access control. (unused on Zen Z0n2p and Zen Z0Hn2p) Precise external termination error (p_tea_b assertion and precise recognition) and $MSR[EE]=1$. See Section 6.2, “Internal Interface Signals,” for a definition of internal signals.
External Input	IVOR 4 ¹	p_extint_b is asserted and $MSR[EE]=1$.

Table 5-2. Exceptions and Conditions (continued)

Interrupt Type	Corresponding Interrupt Vector Offset	Causing Conditions
Alignment	IVOR 5	l mw, stmw not word aligned. lwarx or stwcx . not word aligned.
Program	IVOR 6	Illegal, Privileged, Trap, Unimplemented Operation.
Floating-point unavailable	IVOR 7	Unused
System call	IVOR 8	Execution of the System Call (se_sc) instruction
AP unavailable	IVOR 9	Unused
Decrementer	IVOR 10	Unused
Fixed Interval Timer	IVOR 11	Unused
Watchdog Timer	IVOR 12	Unused
Data TLB Error	IVOR 13	Unused
Instruction TLB Error	IVOR 14	Unused
Debug	IVOR 15	Trap, Instruction Address Compare, Data Address Compare, Instruction Complete, Branch Taken, Return from Interrupt, Interrupt Taken, External Debug Event, Unconditional Debug Event
Reserved	IVOR 16-31	—

¹ Autovector External and Critical Input interrupts use this IVOR. Vectored interrupts supply an interrupt vector offset directly.

5.2 Exception Syndrome Register

The Exception Syndrome Register (ESR) provides a *syndrome* to differentiate between exceptions that can generate the same interrupt type. e200 adds some implementation specific bits to this register, as seen in Figure 5-1.




Figure 5-1. Exception Syndrome Register (ESR)

The ESR bits are defined in [Table 5-3](#).

Table 5-3. ESR Bit Settings

Bit(s)	Name	Description	Associated Interrupt Type
0:3 (32:35)	—	Allocated ¹	—
4 (36)	PIL	Illegal Instruction exception	Program
5 (37)	PPR	Privileged Instruction exception	Program
6 (38)	PTR	Trap exception	Program
7 (39)	FP ²	Floating-point operation	—
8 (40)	ST	Store operation	Alignment Data Storage
9 (41)	—	Reserved ³	—
10 (42)	DLK ⁴	Data Cache Locking	Data Storage
11 (43)	ILK ³	Instruction Cache Locking	Data Storage
12 (44)	AP	Auxiliary Processor operation (Not used by e200)	Alignment (not on e200) Data Storage (not on e200) Data TLB (not on e200) Program (not on e200)
13 (45)	PUO	Unimplemented Operation exception	Program
14 (46)	BO	Byte Ordering exception Mismatched Instruction Storage exception	—
15 (47)	PIE	Program Imprecise exception (Reserved)	Currently unused by e200
16:23 (48:55)	—	Reserved ³	—
24 (56)	—	Reserved	Allocated, is not set by hardware
25 (57)	—	Allocated ¹	—
26 (58)	VLEMI	VLE Mode Instruction	Data Storage Instruction Storage Alignment Program System Call
27:29 (59:61)	—	Allocated ¹	—

Table 5-3. ESR Bit Settings (continued)

Bit(s)	Name	Description	Associated Interrupt Type
30 (62)	MIF	Misaligned Instruction Fetch	Instruction Storage Instruction TLB
31 (63)	XTE	External Termination Error (Precise)	Data Storage Instruction Storage

¹ These bits are not implemented and should be written with zero for future compatibility.

² Unused

³ These bits are not implemented, and should be written with zero for future compatibility.

⁴ Unused

5.3 Machine State Register

The Machine State Register defines the state of the processor. The e200 MSR is shown in Figure 5-2.




Figure 5-2. Machine State Register (MSR)

The MSR bits are defined in Table 5-4.

Table 5-4. MSR Bit Settings

Bit(s)	Name	Description
0:4 (32:36)	—	Reserved ¹
5 (37)	UCLE ²	User Cache Lock Enable 0 Execution of the cache locking instructions in user mode (MSR[PR]=1) disabled; DSI exception taken instead, and ILK or DLK set in ESR. 1 Execution of the cache lock instructions in user mode enabled.
6 (38)	—	Reserved ¹
7:12 (39:44)	—	Reserved ¹

Table 5-4. MSR Bit Settings (continued)

Bit(s)	Name	Description
13 (45)	WE	Wait State (Power management) enable. This bit is defined as optional in the Power Architecture Book E architecture. 0 Power management is disabled. 1 Power management is enabled. The processor can enter a power-saving mode when additional conditions are present. The mode chosen is determined by the DOZE, NAP, and SLEEP bits in the HID0 register, described in Section 2.3.9, “Hardware Implementation Dependent Register 0 (HID0).”
14 (46)	CE	Critical Interrupt Enable 0 Critical Input interrupts are disabled. 1 Critical Input interrupts are enabled.
15 (47)	—	Reserved ¹
16 (48)	EE	External Interrupt Enable 0 External Input interrupts are disabled. 1 External Input interrupts are enabled.
17 (49)	PR	Problem State 0 The processor is in supervisor mode, can execute any instruction, and can access any resource (for example, GPRs, SPRs, MSR, etc.). 1 The processor is in user mode, cannot execute any privileged instruction, and cannot access any privileged resource.
19 (51)	ME	Machine Check Enable 0 Machine Check interrupts are disabled. Checkstop mode is entered when the <i>p_mcp_b</i> input is recognized asserted or an ISI exception occurs on a fetch of the first instruction of an exception handler. 1 Machine Check interrupts are enabled.
20 (52)	FE0	Floating-point exception mode 0 (not used by e200)
21 (53)	—	Reserved ¹
22 (54)	DE	Debug Interrupt Enable 0 Debug interrupts are disabled. 1 Debug interrupts are enabled.
23 (55)	FE1	Floating-point exception mode 1 (not used by e200)
24 (56)	—	Reserved ¹
25 (57)	—	Reserved ¹
26 (58)	IS	Instruction Address Space 0 The processor directs all instruction fetches to address space 0. 1 The processor directs all instruction fetches to address space 1.
27 (59)	DS	Data Address Space 0 The processor directs all data storage accesses to address space 0. 1 The processor directs all data storage accesses to address space 1.

Table 5-4. MSR Bit Settings (continued)

Bit(s)	Name	Description
28:29 (60:61)	—	Reserved ¹
30 (62)	RI	Recoverable Interrupt 0 Machine Check interrupt is not recoverable. 1 Machine Check interrupt may be recoverable. This bit is cleared when a Machine check or critical class interrupt which uses CSRR0/1 is taken (see Table 5-1). It is not set by hardware, and does not affect processor operation. It is provided as a software assist to determine if machine check interrupts may possibly be recoverable.
30:31 (62:63)	—	Reserved ¹

¹ These bits are not implemented, are read as zero, and writes are ignored.

² This bit is implemented, but ignored

5.4 Machine Check Syndrome Register (MCSR)

When the core complex takes a machine check interrupt, it updates the Machine Check Syndrome register (MCSR) to differentiate between machine check conditions. The MCSR is shown in Figure 5-3.




Figure 5-3. Machine Check Syndrome Register (MCSR)

Table 5-5 describes MCSR fields. The MCSR indicates the source of a machine check condition is recoverable. When a syndrome bit in the MCSR is set, the core complex asserts *p_mcp_out* for system information.

Table 5-5. Machine Check Syndrome Register (MCSR)

Bit	Name	Description	Recoverable
0 (32)	MCP	Machine check input pin	Maybe
1 (33)	—	Reserved, should be cleared.	—

Table 5-5. Machine Check Syndrome Register (MCSR) (continued)

Bit	Name	Description	Recoverable
2 (34)	CP_PERR ¹	Cache push parity error	Unlikely
3 (35)	CPERR ¹	Cache parity error	Precise
4 (36)	EXCP_ERR	Bus Error on first instruction fetch for an exception handler	Precise
e200z0: 5:26 (37: 58)) e200z0h: 5:10(37: 42)	—	Reserved, should be cleared.	—
11 (43)	NMI	Non-maskable interrupt input pin	Maybe
12:26 (44:58)	—	Reserved, should be cleared.	—
27 (59)	BUS_IRERR	Read bus error on Instruction fetch	Unlikely
28 (60)	BUS_DRERR	Read bus error on data load	Unlikely
29 (61)	BUS_WRERR	Write bus error on data store	Unlikely
30:31 (62:63)	—	Reserved, should be cleared.	—

¹ This bit is implemented, but is never set by hardware.

5.5 Interrupt Vector Prefix Register (IVPR)

The Interrupt Vector Prefix Register is used during interrupt processing for determining the starting address of a software handler used to handle an interrupt. The hardwired Interrupt Vector Offset value for a particular interrupt type is concatenated with the value held in the Interrupt Vector Prefix register (IVPR) to form an instruction address from which execution is to begin. Note that for Zen Z0n2p and Zen Z0Hn2p the IVPR has been extended from 16 to 20 bits, allowing the vector table to reside on any 4-Kbyte boundary. The format of IVPR is shown in [Figure 5-4](#).




Figure 5-4. e200 Interrupt Vector Prefix Register (IVPR)

The IVPR fields are defined in [Table 5-6](#).

Table 5-6. IVPR Register Fields

Bit(s)	Name	Description
0:19 (32:51)	Vec Base	Vector Base This field is used to define the base location of the vector table, aligned to a 4Kbyte boundary. This field provides the high-order 20 bits of the location of all interrupt handlers. The IVORxx value appropriate for the type of exception being processed are concatenated with the IVPR Vector Base to form the address of the handler in memory.
20:31 (52:63)	—	Reserved ¹

¹ These bits are not implemented, are read as zero, and writes are ignored.

5.6 Interrupt Vector Offset Values (IVORxx)

Interrupt Vector Offset Registers are not implemented in Zen Z0n2p and Zen Z0Hn2p. Instead, hardwired values are used for interrupt offsets to the IVPR contents during interrupt processing for determining the starting address of a software handler used to handle an interrupt. The associated hardwired value of the Interrupt Vector Offset selected for a particular interrupt type is concatenated with the value held in the Interrupt Vector Prefix register (IVPR) to form an instruction address from which execution is to begin as shown in [Figure 5-5](#).




Figure 5-5. e200 Interrupt Vector Addresses

The hardwired Vector Offsets are listed in [Table 5-7](#).

Table 5-7. Hardwired Vector Offset Values

Interrupt Type	Corresponding Interrupt Vector Offset Register	12-Bit Hex Offset
Critical Input	IVOR 0 ¹	0x000
Machine check	IVOR 1	0x010
Data Storage	IVOR 2	0x020
Instruction Storage	IVOR 3	0x030
External Input	IVOR 4 ¹	0x040
Alignment	IVOR 5	0x050
Program	IVOR 6	0x060
Floating-point unavailable	IVOR 7	Unused
System call	IVOR 8	0x080
AP unavailable	IVOR 9	Unused
Decrementer	IVOR 10	Unused

Table 5-7. Hardwired Vector Offset Values (continued)

Interrupt Type	Corresponding Interrupt Vector Offset Register	12-Bit Hex Offset
Fixed Interval Timer	IVOR 11	Unused
Watchdog Timer	IVOR 12	Unused
Data TLB Error	IVOR 13	Unused
Instruction TLB Error	IVOR 14	Unused
Debug	IVOR 15	0x0F0

¹ Autovectored External and Critical Input interrupts use this IVOR. Vectored interrupts supply an interrupt vector offset directly.

5.7 Interrupt Definitions

5.7.1 Critical Input Interrupt (IVOR0)

A Critical Input exception is signalled to the processor by the assertion of the critical interrupt pin (*p_critint_b*). When e200 detects the exception, if the exception is enabled by MSR[CE], e200 takes the Critical Input interrupt. The *p_critint_b* input is a level-sensitive signal expected to remain asserted until e200 acknowledges the interrupt. If *p_critint_b* is negated early, recognition of the interrupt request is not guaranteed. After e200 begins execution of the critical interrupt handler, the system can safely negate *p_critint_b*.

A Critical Input interrupt may be delayed by other higher priority exceptions or if MSR[CE] is cleared when the exception occurs.

Table 5-8 lists register settings when a Critical Input interrupt is taken.

Table 5-8. Critical Input Interrupt—Register Settings

Register	Setting Description		
CSRR0	Set to the effective address of the instruction that the processor would have attempted to execute next if no exception conditions were present.		
CSRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE 0 EE 0 PR 0	FP 0 ME — FE0 0 DE —/0 ¹	FE1 0 IS 0 DS 0 RI 0 ²
ESR	Unchanged		
MCSR	Unchanged		
DEAR	Unchanged		
Vector	IVPR _{0:19} 12'h000 (autovectored) IVPR _{0:19} <i>p_voffset</i> [0:9] 2'b00 (non-autovectored)		

- ¹ DE is cleared when the Debug APU is disabled. Clearing of DE is optionally supported by control in HID0 when the Debug APU is enabled.
- ² RI is cleared by all critical class interrupts using CSRR0/1 and the machine check interrupt. These interrupt handlers should set RI to '1' early in the handler after CSRR0/1 have been saved to allow for improved recoverability.

When the Debug APU is enabled, the MSR[DE] bit is not automatically cleared by a Critical Input interrupt, but can be configured to be cleared via the HID0 register (HID0[CICLRDE]). Refer to [Section 2.3.9, “Hardware Implementation Dependent Register 0 \(HID0\).”](#)

IVOR0 is the vector offset used by autovector Critical Input interrupts to determine the interrupt handler location. e200 also provides the capability to directly vector Critical Input interrupts to multiple handlers by allowing a Critical Input interrupt request to be accompanied by a vector offset. The *p_voffset*[0:9] input signals are appended with 2'b00 and used in place of the IVOR0 value to form the interrupt vector when a Critical Input interrupt request is not autovector (*p_avec_b* negated when *p_critint_b* asserted).

5.7.2 Machine Check Interrupt (IVOR1)

e200 implements the Machine Check exception as defined in Power Architecture Book E except for automatic clearing of the MSR[DE] bit (see later paragraph). e200 initiates a Machine Check interrupt if MSR[ME]=1 and any of the machine check sources listed in [Table 5-2](#) is detected. As defined in Power Architecture Book E, the interrupt is not taken if MSR[ME] is cleared, in which case the processor generates an internal checkstop condition and enters the checkstop state. When a processor is in the checkstop state, instruction processing is suspended and generally cannot continue without restarting the processor. Note that other conditions may lead to the checkstop condition; the disabled machine check exception is only one of these.

e200 implements the Machine Check Syndrome register (MCSR) to record the source(s) of machine checks.

The MSR[DE] bit is not automatically cleared by a Machine Check exception, but can be configured to be cleared or left unchanged via the HID0 register (HID0[MCCLRDE]). Refer to [Section 2.3.9, “Hardware Implementation Dependent Register 0 \(HID0\).”](#)

5.7.2.1 Machine Check Interrupt Enabled (MSR[ME]=1)

Machine Check interrupts are enabled when MSR[ME]=1. When a Machine Check interrupt is taken, registers are updated as shown in [Table 5-9](#).

Table 5-9. Machine Check Interrupt—Register Settings

Register	Setting Description
CSRR0	On a best-effort basis e200 sets this to the address of some instruction that was executing or about to be executing when the machine check condition occurred.
CSRR1	Set to the contents of the MSR at the time of the interrupt

Table 5-9. Machine Check Interrupt—Register Settings (continued)

Register	Setting Description		
MSR	UCLE 0 WE 0 CE 0 EE 0 PR 0	FP 0 ME 0 FE0 0 DE 0/— ¹	FE1 0 IS 0 DS 0 RI 0 ²
ESR	Unchanged		
MCSR	Updated to reflect the source(s) of a machine check		
DEAR	Unchanged		
Vector	IVPR _{0:19} 12'h010		

¹ DE is cleared when the Debug APU is disabled. Clearing of DE is optionally supported by control in HID0 when the Debug APU is enabled.

² RI is cleared by all critical class interrupts using CSRR0/1 and the machine check interrupt. These interrupt handlers should set RI to '1' early in the handler after CSRR0/1 have been saved to allow for improved recoverability.

The Machine Check input pin *p_mcp_b* can be masked by HID0[EMCP].

In general, most Machine Check exceptions are unrecoverable in the sense that execution cannot resume in the context that existed before the interrupt; however, system software can use the machine check interrupt handler to try to identify and recover from the machine check condition.

The Machine Check Syndrome register is provided to identify the source(s) of a machine check and may be used to identify recoverable events.

The interrupt handler should set MSR[ME] as soon as possible to avoid entering the checkstop state if another machine check condition were to occur.

5.7.2.2 Checkstop State

A checkstop condition can occur for several reasons. The exception related conditions are:

- MSR[ME]=0 and a machine check occurs (other than a non-maskable interrupt on e200z0h).
- First instruction in an interrupt handler can not be executed due to a bus error termination, and MSR[ME]=0.
- Bus error termination for a buffered store and MSR[ME]=0.
- Precise external termination error and MSR[EE]=0 and MSR[ME]=0

When a processor is in the checkstop state, instruction processing is suspended and generally cannot resume without the processor being reset. To indicate that a checkstop condition exists, the *p_chkstop* output pin is asserted whenever the CPU is in the checkstop state.

When a debug request is presented to the e200 core while in the checkstop state, the *p_wakeup* signal is asserted, and when *m_clk* is provided to the CPU, it temporarily exits the checkstop state and enters Debug mode. The *p_chkstop* output is negated for the duration of the time the CPU remains in a debug session (*p_debug_b* asserted). When the debug session is exited, the CPU re-enters the checkstop state. Note that the external system logic may be in an undefined state following a checkstop condition, such as having an

outstanding bus transaction, or other inconsistency, thus no guarantee can be made in general about activities performed in debug mode while a checkstop is still outstanding. Debug logic does have the capability of generating assertion of the *p_resetout_b* signal via the DBCR0 register though.

5.7.2.3 Non-Maskable Interrupts (NMI) in e200z0h

e200 implements a non-maskable interrupt in addition to the machine check sources defined by Power Architecture Book E. The non-maskable interrupt is signaled via the *p_nmi_b* input. Non-maskable interrupt are *not* gated by MSR[ME], and a non-maskable interrupt occurring with MSR[ME]=0 does *not* result in a checkstop condition. e200 provides the MSR[RI] bit to indicate whether these non-maskable interrupts are potentially recoverable. Because a non-maskable interrupt overwrites the CSRR0/1 registers, if these registers are currently holding essential state because a critical class interrupt handler has not yet been able to save this state away safely, and a non-maskable interrupt occurs, no recovery from the earlier critical class interrupt is possible. The machine check handler may use the value of CSRR1[RI] to determine if this has occurred. If CSRR1[RI] is cleared, then no recovery is possible, because MSR[RI] was 0 at the time of the non-maskable interrupt, indicating that the CSRR0/1 registers had not yet been saved. Critical class and machine check interrupt handlers should save the state of CSRR0/1 and then set MSR[RI] to '1' as soon as is practical to ensure the best chance of recovery from a non-maskable interrupt.

5.7.3 Data Storage Interrupt (IVOR2)

A Data Storage interrupt (DSI) may occur if no higher priority exception exists and one of the following exception conditions exists:

- External Termination Error (precise) and MSR[EE]=1

Precise external termination errors occur when a load or store is terminated by assertion of *p_tea_b* (ERROR termination response). See [Section 6.2, “Internal Interface Signals,”](#) for a definition of internal signals.

[Table 5-10](#) lists register settings when a DSI is taken.

Table 5-10. Data Storage Interrupt—Register Settings

Register	Setting Description		
SRR0	Set to the effective address of the excepting load/store instruction.		
SRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE — EE 0 PR 0	FP 0 ME — FE0 0 DE —	FE1 0 IS 0 DS 0 RI —
ESR	Access: External Termination Error (Precise):	[ST], [VLEMI]. All other bits cleared. [ST], [VLEMI], XTE. All other bits cleared.	
MCSR	Unchanged		
DEAR	For Access exceptions, set to the effective address of a byte within the page whose access caused the violation.		
Vector	IVPR _{0:19} 12'h020		

5.7.4 Instruction Storage Interrupt (IVOR3)

An Instruction Storage interrupt (ISI) occurs when no higher priority exception exists and a precise external termination error occurs when an instruction fetch is terminated by assertion of *p_tea_b* (ERROR termination response) and MSR[EE]=1. See [Section 6.2, “Internal Interface Signals,”](#) for a definition of internal signals.

[Table 5-11](#) lists register settings when an ISI is taken.

Table 5-11. Instruction Storage Interrupt—Register Settings

Register	Setting Description		
SRR0	Set to the effective address of the excepting instruction.		
SRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE — EE 0 PR 0	FP 0 ME — FE0 0 DE —	FE1 0 IS 0 DS 0 RI —
ESR	XTE, VLEMI. All other bits cleared.		
MCSR	Unchanged		
DEAR	Unchanged		
Vector	IVPR _{0:19} 12'h030		

5.7.5 External Input Interrupt (IVOR4)

An External Input exception is signalled to the processor by the assertion of the external interrupt pin (*p_extint_b*). The *p_extint_b* input is a level-sensitive signal expected to remain asserted until e200 acknowledges the external interrupt. If *p_extint_b* is negated early, recognition of the interrupt request is not guaranteed. When e200 detects the exception, if the exception is enabled by MSR[EE], e200 takes the External Input interrupt.

An External Input interrupt may be delayed by other higher priority exceptions or if MSR[EE] is cleared when the exception occurs.

[Table 5-12](#) lists register settings when an External Input interrupt is taken.

Table 5-12. External Input Interrupt—Register Settings

Register	Setting Description		
SRR0	Set to the effective address of the instruction that the processor would have attempted to execute next if no exception conditions were present.		
SRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE — EE 0 PR 0	FP 0 ME — FE0 0 DE —	FE1 0 IS 0 DS 0 RI —

Table 5-12. External Input Interrupt—Register Settings (continued)

ESR	Unchanged
MCSR	Unchanged
DEAR	Unchanged
Vector	IVPR _{0:19} 12'h040 (autovectored) IVPR _{0:19} <i>p_voffset</i> [0:9] 2'b00 (non-autovectored)

IVOR4 is the vector offset value used by autovectored External Input interrupts to determine the interrupt handler location. e200 also provides the capability to directly vector External Input interrupts to multiple handlers by allowing a External Input interrupt request to be accompanied by a vector offset. The *p_voffset*[0:9] input signals are appended with 2'b00 and used in place of the IVOR4 value when an External Input interrupt request is not autovectored (*p_avec_b* negated when *p_extint_b* asserted).

5.7.6 Alignment Interrupt (IVOR5)

e200 implements the Alignment Interrupt as defined by Power Architecture Book E. An Alignment exception is generated when any of the following occurs:

- The operand of **e_lmw** or **e_stmw** not word aligned
- The operand of **lwarx** or **stwcx.** not word aligned

Execution of a **dcbz** instruction is attempted [Table 5-13](#) lists register settings when an alignment interrupt is taken.

Table 5-13. Alignment Interrupt—Register Settings

Register	Setting Description		
SRR0	Set to the effective address of the excepting load/store instruction.		
SRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE — EE 0 PR 0	FP 0 ME — FE0 0 DE —	FE1 0 IS 0 DS 0 RI —
ESR	[ST], VLEMI. All other bits cleared.		
MCSR	Unchanged		
DEAR	Set to the effective address of a byte of the load or store whose access caused the violation.		
Vector	IVPR _{0:19} 12'h050		

5.7.7 Program Interrupt (IVOR6)

e200 implements the Program Interrupt as defined by Power Architecture Book E. A program interrupt occurs when no higher priority exception exists and one or more of the following exception conditions defined in Power Architecture Book E occur:

- Illegal Instruction exception
- Privileged Instruction exception
- Trap exception
- Unimplemented Operation exception

e200 invokes an Illegal Instruction program exception on attempted execution of the following instructions:

- Instruction from the illegal instruction class
- **mtspr** and **mfspir** instructions with an undefined SPR specified
- **mtdcr** and **mfddcr** instructions with an undefined DCR specified

e200 invokes a Privileged Instruction program exception on attempted execution of the following instructions when MSR[PR]=1 (user mode):

- A privileged instruction
- **mtspr** and **mfspir** instructions which specify a SPRN value with SPRN₅=1 (even if the SPR is undefined).

e200 invokes an Trap exception on execution of **tw** instructions if the trap conditions are met and the exception is not also enabled as a Debug interrupt.

All other defined or allocated instructions that are not implemented by e200 cause a illegal instruction program exception.

Table 5-14 lists register settings when a Program interrupt is taken.

Table 5-14. Program Interrupt—Register Settings

Register	Setting Description		
SRR0	Set to the effective address of the excepting instruction.		
SRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE — EE 0 PR 0	FP 0 ME — FE0 0 DE —	FE1 0 IS 0 DS 0 RI —
ESR	Illegal: Privileged: Trap: Unimplemented:	PIL, VLEMI. All other bits cleared. PPR, VLEMI. All other bits cleared. PTR, VLEMI. All other bits cleared. PUO, VLEMI. All other bits cleared.	
MCSR	Unchanged		
DEAR	Unchanged		
Vector	IVPR _{0:19} 12'h060		

5.7.8 System Call Interrupt (IVOR8)

A System Call interrupt occurs when a System Call (**se_sc**) instruction is executed and no higher priority exception exists.

Exception extensions implemented in e200 for Power Architecture VLE include modification of the System Call Interrupt definition to include updating the ESR.

Table 5-15 lists register settings when a System Call interrupt is taken.

Table 5-15. System Call Interrupt—Register Settings

Register	Setting Description		
SRR0	Set to the effective address of the instruction <i>following</i> the sc instruction.		
SRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE — EE 0 PR 0	FP 0 ME — FE0 0 DE —	FE1 0 IS 0 DS 0 RI —
ESR	VLEMI All other bits cleared.		
MCSR	Unchanged		
DEAR	Unchanged		
Vector	IVPR _{0:19} 12'h080		

5.7.9 Auxiliary Processor Unavailable Interrupt (IVOR9)

An Auxiliary Processor Unavailable exception is defined by Power Architecture Book E to occur when an attempt is made to execute an APU instruction which is implemented but configured as unavailable, and no higher priority exception condition exists.

e200 does not utilize this interrupt.

5.7.10 Debug Interrupt (IVOR15)

e200 implements the Debug Interrupt as defined in Power Architecture Book E with the following changes:

- When the Debug APU is enabled, Debug is no longer a critical interrupt, but uses DSRR0 and DSRR1 for saving machine state on context switch
- A Return from debug interrupt instruction (**se_rfdi**) is implemented to support the new machine state registers
- A Critical Interrupt Taken debug event is defined to allow critical interrupts to generate a debug event
- A Critical Return debug event is defined to allow debug events to be generated for **se_rfdi** instructions

There are multiple sources that can signal a Debug exception. A Debug interrupt occurs when no higher priority exception exists, a Debug exception exists in the Debug Status Register, and Debug interrupts are enabled (both DBCR0[IDM]=1 (internal debug mode) and MSR[DE]=1). Enabling debug events and other debug modes are discussed further in [Chapter 8, “Debug Support.”](#) With the Debug APU enabled, (See [Section 2.3.9, “Hardware Implementation Dependent Register 0 \(HID0\)”](#)) the Debug interrupt has its own set of machine state save/restore registers (DSRR0, DSRR1) to allow debugging of both critical and non-critical interrupt handlers. In addition, the capability is provided to allow interrupts to be handled while in a debug software handler. External and Critical interrupts are not automatically disabled when a Debug interrupt occurs but can be configured to be cleared via the HID0 register (HID0[DCLREE, DCLRCE]). Refer to [Section 2.3.9, “Hardware Implementation Dependent Register 0 \(HID0\).”](#) When the Debug APU is disabled, Debug interrupts use the CSRR0 and CSRR1 registers to save machine state.

An Instruction Address Compare (IAC) debug exception occurs when there is an instruction address match as defined by the debug control registers and Instruction Address Compare events are enabled. This could either be a direct instruction address match or a selected set of instruction addresses. IAC has the highest interrupt priority of all instruction-based interrupts, even if the instruction itself may have encountered an Instruction Storage exception.

A Branch Taken (BRT) debug exception is signalled when a branch instruction is considered taken by the branch unit and branch taken events are enabled. The Debug interrupt is taken when no higher priority exception is pending.

A Data Address Compare (DAC) exception is signalled when there is a data access address match as defined by the debug control registers and Data Address Compare events are enabled. This could either be a direct data address match or a selected set of data addresses. The Debug interrupt is taken when no higher priority exception is pending.

e200 does not implement the Data Value Compare debug mode, specified in Power Architecture Book E.

The e200 implementation provides IAC linked with DAC exceptions. This results in a DAC exception only if one or more IAC conditions are also met. See [Chapter 8, “Debug Support,”](#) for more details.

A Trap (TRAP) debug exception occurs when a program trap exception is generated while trap events are enabled. If MSR[DE] is set, the Debug exception has higher priority than the Program exception in this case, and is taken instead of a Trap type Program Interrupt. The Debug interrupt is taken when no higher priority exception is pending. If MSR[DE] is cleared when a trap debug exception occurs, a Trap exception type Program interrupt occurs instead.

A Return (RET) debug exception occurs when executing an **se_rfi** instruction and return debug events are enabled. Return debug exceptions are not generated for **se_rfci** or **se_rfdi** instructions. If MSR[DE]=1 at the time of the execution of the **se_rfi**, a Debug interrupt occurs provided there exists no higher priority exception which is enabled to cause an interrupt. CSRR0 (Debug APU disabled) or DSRR0 (Debug APU enabled) is set to the address of the **se_rfi** instruction. If MSR[DE]=0 at the time of the execution of the **se_rfi**, a Debug interrupt does not occur immediately, but the event is recorded by setting the DBSR[RET] and DBSR[IDE] status bits.

A Critical Return (CRET) debug exception occurs when executing an **se_rfci** instruction and critical return debug events are enabled. Critical return debug exceptions are only generated for **se_rfci** instructions. If MSR[DE]=1 at the time of the execution of the **se_rfci**, a Debug interrupt occurs provided there exists no

higher priority exception which is enabled to cause an interrupt. CSRR0 (Debug APU disabled) or DSRR0 (Debug APU enabled) is set to the address of the **se_rfc**i instruction. If MSR[DE]=0 at the time of the execution of the **se_rfc**i, a Debug interrupt does not occur immediately, but the event is recorded by setting the DBSR[CRET] and DBSR[IDE] status bits. Note that critical return debug events should not normally be enabled unless the Debug APU is enabled to avoid corruption of CSRR0/1.

An Instruction Complete (ICMP) debug exception is signalled following execution and completion of an instruction while this event is enabled.

A **mtmsr** or **mtdbcr0** which causes both MSR[DE] and DBCR0[IDM] to end up set, enabling precise debug mode, may cause an Imprecise (Delayed) Debug exception to be generated due to an earlier recorded event in the Debug Status register.

An Interrupt Taken (IRPT) debug exception occurs when a non-critical interrupt context switch is detected. This exception is imprecise and unordered with respect to the program flow. Note that an IRPT Debug interrupt only occurs when detecting a non-critical interrupt on e200. The value saved in CSRR0/DSRR0 is the address of the non-critical interrupt handler.

A Critical Interrupt Taken (CIRPT) debug exception occurs when a critical interrupt context switch is detected. This exception is imprecise and unordered with respect to the program flow. Note that a CIRPT Debug interrupt only occurs when detecting a critical interrupt on e200. The value saved in CSRR0/DSRR0 is the address of the critical interrupt handler. Note that Critical Interrupt Taken debug events should not normally be enabled unless the Debug APU is enabled to avoid corruption of CSRR0/1.

An Unconditional Debug Event (UDE) exception occurs when the Unconditional Debug Event pin (*p_ude*) transitions to the asserted state.

External Debug exceptions occur when enabled and one of the External Debug Event pins (*p_dev1*, *p_dev2*) transitions to the asserted state.

The Debug Status Register (DBSR) provides a *syndrome* to differentiate between debug exceptions that can generate the same interrupt. For more details see [Chapter 8, “Debug Support.”](#)

[Table 5-16](#) lists register settings when a Debug interrupt is taken.

Table 5-16. Debug Interrupt—Register Settings

Register	Setting Description		
CSRR0/ DSRR0 ¹	Set to the effective address of the excepting instruction for IAC, BRT, RET, CRET, and TRAP. Set to the effective address of the next instruction to be executed <i>following</i> the excepting instruction for DAC and ICMP. For a UDE, IRPT, CIRPT, DCNT, or DEVT type exception, set to the effective address of the instruction that the processor would have attempted to execute next if no exception conditions were present.		
CSRR1/ DSRR1	Set to the contents of the MSR at the time of the interrupt		
MSR	UCLE 0 WE 0 CE —/0 ² EE —/0 ² PR 0	FP 0 ME — FE0 0 DE 0	FE1 0 IS 0 DS 0 RI —/0 ^{2,3}

Table 5-16. Debug Interrupt—Register Settings (continued)

DBSR ⁴	Unconditional Debug Event: Instr. Complete Debug Event: Branch Taken Debug Event: Interrupt Taken Debug Event: Critical Interrupt Taken Debug Event: Trap Instruction Debug Event: Instruction Address Compare: Data Address Compare: Return Debug Event: Critical Return Debug Event: External Debug Event: and optionally, an Imprecise Debug Event flag	UDE ICMP BRT IRPT CIRPT TRAP {IAC1, IAC2, IAC3, IAC4} {DAC1R, DAC1W, DAC2R, DAC2W} RET CRET {DEVT1, DEVT2} {IDE}
ESR	Unchanged	
MCSR	Unchanged	
DEAR	Unchanged	
Vector	IVPR _{0:19} 12'h0F0	

¹ Assumes that the Debug interrupt is precise

² Conditional based on control bits in HID0. If HID0[DAPUEN] = 1, RI is unaffected because DSRR0/1 are used, otherwise it is cleared because CSRR0/1 are updated.

³ RI is cleared by all critical class interrupts using CSRR0/1 and the machine check interrupt. These interrupt handlers should set RI to '1' early in the handler after CSRR0/1 have been saved to allow for improved recoverability.

⁴ Note that multiple DBSR bits may be set

5.7.11 System Reset Interrupt

e200 implements the System Reset interrupt as defined in Power Architecture Book E. The System Reset exception is a non-maskable, asynchronous exception signalled to the processor through the assertion of system-defined signals.

A System reset may be initiated by either asserting the *p_reset_b* input signal, during power-on reset by asserting *m_por*, by Watchdog Timer Reset Control, or by Debug Reset Control. The *m_por* signal must be asserted during power up and must remain asserted for a period that allows internal logic to be reset.

The *p_reset_b* signal must also remain asserted for a period that allows internal logic to be reset. This period is specified in the hardware specifications. If *m_por* or *p_reset_b* are asserted for less than the required interval, the results are not predictable.

When a reset request occurs, the processor branches to the system reset exception vector (value on *p_rstbase*[0:29] concatenated with 2'b00) without attempting to reach a recoverable state. If reset occurs during normal operation, all operations cease and the machine state is lost. e200 internal state after a reset is defined in [Section 2.4.4, “Reset Settings.”](#)

For reset initiated by Debug Reset Control, e200 implements DBSR[MRR] to aid software in determining the cause. Debug Reset Control provides the capability to assert the *p_resetout_b* signal. External logic may factor this signal into the *p_reset_b* input signal to cause a e200 reset to occur.

Table 5-17 shows the DBSR register bits associated with reset status.

Table 5-17. DBSR Most Recent Reset

Bit(s)	Name	Function
2:3 (34:35)	MRR	00 No reset occurred because these bits were last cleared by software 01 A reset occurred because these bits were last cleared by software 10 Reserved 11 Reserved

Table 5-18 lists register settings when a System Reset interrupt is taken.

Table 5-18. System Reset Interrupt—Register Settings

Register	Setting Description		
CSRR0	Undefined.		
CSRR1	Undefined.		
MSR	UCLE 0 WE 0 CE 0 EE 0 PR 0	FP 0 ME 0 FE0 0 DE 0	FE1 0 IS 0 DS 0 RI 0
ESR	Cleared		
DEAR	Undefined		
Vector	[<i>p_rstbase</i> [0:29]] 2'b00		

5.8 Exception Recognition and Priorities

The following list of exception categories describes how e200 handles exceptions up to the point of signaling the appropriate interrupt to occur. Also, instruction completion is defined as updating all architectural registers associated with that instruction as necessary, and then removing the instruction from the pipeline.

- Interrupts caused by asynchronous events (exceptions). These exceptions are further distinguished by whether they are maskable and recoverable.
 - Asynchronous, non-maskable, non-recoverable:
System reset by assertion of *p_reset_b*
Has highest priority and is taken immediately regardless of other pending exceptions or recoverability. (Includes Debug Reset Control)
 - Asynchronous, non-maskable, possibly non-recoverable:
Non-maskable interrupt by assertion of *p_nmi_b*
Has priority over any other pending exception except system reset conditions.
Recoverability is dependent on whether CSRR0/1 are holding essential state info and are overwritten when the NMI occurs.
 - Asynchronous, maskable, possibly non-recoverable:

Machine check interrupt

Has priority over any other pending exception except system reset conditions.

Recoverability is dependent on the source of the exception. Typically unrecoverable.

— Asynchronous, maskable, recoverable:

External Input, Critical Input, Unconditional Debug, and External Debug Event interrupts

Before handling this type of exception, the processor needs to reach a recoverable state. A maskable recoverable exception remains pending until taken or cancelled by software.

- Synchronous, non instruction-based interrupts. The only exception in this category is the Interrupt Taken debug exception, recognized by an interrupt taken event. It is not considered instruction-based but is synchronous with respect to the program flow.

— Synchronous, maskable, recoverable:

Interrupt Taken debug event.

The machine is in a recoverable state due to the state of the machine at the context switch triggering this event.

- Instruction-based interrupts. These interrupts are further organized by the point in instruction processing in which they generate an exception.

— Instruction Fetch:

Instruction Storage and Instruction Address Compare debug exceptions.

Once these types of exceptions are detected, the excepting instruction is tagged. When the excepting instruction is next to begin execution and a recoverable state has been reached, the interrupt is taken. If an event prior to the excepting instruction causes a redirection of execution, the instruction fetch exception is discarded (but may be encountered again).

— Instruction Dispatch/Execution:

Program, System Call, Data Storage, Alignment, Debug (Trap, Branch Taken, Ret) interrupts.

These types of exceptions are determined during decode or execution of an instruction. The exception remains pending until all instructions before the exception causing instruction in program order complete. The interrupt is then taken without completing the exception-causing instruction. If completing previous instructions causes an exception, that exception takes priority over the pending instruction dispatch/execution exception, which is discarded (but may be encountered again when instruction processing resumes).

— Post-Instruction Execution

Debug (Data Address Compare, Instruction Complete) interrupt.

These Debug exceptions are generated following execution and completion of an instruction while the event is enabled. If executing the instruction produces conditions for another type of exception with higher priority, that exception is taken and the post-instruction exception is discarded for the instruction (but may be encountered again when instruction processing resumes)

5.8.1 Exception Priorities

Exceptions are prioritized as described in [Table 5-19](#). Some exceptions may be masked or imprecise, which affect their priority. Non-maskable exceptions such as reset and machine check may occur at any time and are not delayed even if an interrupt is being serviced, thus state information for any interrupt may be lost. Reset and most machine checks are non-recoverable.

Table 5-19. e200 Exception Priorities

Priority	Exception	Cause	IVOR
Asynchronous Exceptions			
0	System reset	Assertion of <i>p_reset_b</i> or Debug Reset Control	none
1	Machine check	Assertion of <i>p_mcp_b</i> , exception on fetch of first instruction of an interrupt handler, bus error on buffered store, Bus error (XTE) with MSR[EE]=0 and current MSR[ME]=1 ³ , assertion of <i>p_nmi_b</i>	1
2	—	—	—
3 ¹	Debug: 1. UDE 2. DEVT1 3. DEVT2 4. DCNT1 ² 5. DCNT2 ² 6. IDE	1. Assertion of <i>p_ude</i> (Unconditional Debug Event) 2. Assertion of <i>p_devt1</i> and event enabled (External Debug Event 1) 3. Assertion of <i>p_devt2</i> and event enabled (External Debug Event 2) 4. Debug Counter 1 exception 5. Debug Counter 2 exception 6. Imprecise Debug Event (event imprecise due to previous higher priority interrupt)	15
4 ¹	Critical Input	Assertion of <i>p_critint_b</i>	0
5 ^{1,2}	Watchdog Timer	Watchdog Timer first enabled time-out	12
6 ¹	External Input	Assertion of <i>p_extint_b</i>	4
7 ^{1,2}	Fixed-Interval Timer	Posting of a FIT exception in TSR due to programmer-specified bit transition in the Time Base register	11
8 ^{1,2}	Decrementer	Posting of a Decrementer exception in TSR due to programmer-specified Decrementer condition	10
Instruction Fetch Exceptions			
9	Debug: IAC (unlinked)	Instruction address compare match for enabled IAC debug event and DBCR0[IDM] asserted	15
10 ²	ITLB Error	Instruction translation lookup miss in the TLB	14
11	Instruction Storage	1. Access control. (unused on Zen Z0n2pZen Z0Hn2p) 2. Precise external termination error (<i>p_tea_b</i> assertion and precise recognition) and MSR[EE]=1	3
Instruction Dispatch/Execution Interrupts			
12	Program: Illegal	Attempted execution of an illegal instruction.	6
13	Program: Privileged	Attempted execution of a privileged instruction in user-mode	6

Table 5-19. e200 Exception Priorities (continued)

Priority	Exception	Cause	IVOR
14 ²	Floating-point Unavailable	Any floating-point unavailable exception condition.	7
15	Program: Unimplemented	Attempted execution of an unimplemented instruction.	6
16	Debug: 1. BRT 2. Trap 3. RET 4. CRET	1. Attempted execution of a taken branch instruction 2. Condition specified in tw instruction met. 3. Attempted execution of a se_rfi instruction. 4. Attempted execution of an se_rfc instruction. Note: Exceptions requires corresponding debug event enabled, MSR[DE]=1, and DBCR0[IDM]=1.	15
17	Program: Trap	Condition specified in tw instruction met and not trap debug.	15
	System Call	Execution of the System Call (se_sc) instruction.	8
18	Alignment	lmw , stmw , lwarx , or stwcx . not word aligned. dcbz with cache disabled or not present	5
19	Debug with concurrent DSI exception: 1. DAC/IAC linked ³ 2. DAC unlinked ³	Debug with concurrent DSI exception. DBSR[IDE] also set. Data Address Compare linked with Instruction Address Compare Data Address Compare unlinked Note: Exceptions requires corresponding debug event enabled, MSR[DE]=1, and DBCR0[IDM]=1. In this case, the Debug exception is considered imprecise, and DBSR[IDE] is set. Saved PC points to the load or store instruction causing the DAC event.	15
20 ²	Data TLB Error	Data translation lookup miss in the TLB.	13
21	Data Storage	1. Access control. (unused on Zen Z0n2pZen Z0Hn2p) 2. Precise external termination error (<i>p_tea_b</i> assertion and precise recognition) and MSR[EE]=1	2
22 ^{2,4}	Alignment	dcbz to W=1 or I=1 storage with cache enabled	5
23	Debug: 1. IRPT 2. CIRPT	1. Interrupt taken (non-critical) 2. Critical Interrupt taken (critical only) Note: Exceptions requires corresponding debug event enabled, MSR[DE]=1, and DBCR0[IDM]=1.	15

Table 5-19. e200 Exception Priorities (continued)

Priority	Exception	Cause	IVOR
Post-Instruction Execution Exceptions			
24	Debug: 1. DAC/IAC linked ³ 2. DAC unlinked ³	1. Data Address Compare linked with Instruction Address Compare 2. Data Address Compare unlinked Notes: Exceptions requires corresponding debug event enabled, MSR[DE]=1, and DBCR0[IDM]=1. Saved PC points to the instruction following the load or store instruction causing the DAC event.	15
25	Debug: 1. ICMP	1. Completion of an instruction. Note: Exceptions requires corresponding debug event enabled, MSR[DE]=1, and DBCR0[IDM]=1.	15

¹ These exceptions are sampled at instruction boundaries, thus may actually occur after exceptions which are due to a currently executing instruction. If one of these exceptions occurs during execution of an instruction in the pipeline, it is not processed until the pipeline has been flushed, and the exception associated with the excepting instruction may occur first.

² Unused on Zen Z0n2p and Zen Z0Hn2p

³ When no Data Storage Interrupt or Data TLB Error occurs, e200 implements the data address compare debug exceptions as post-instruction exceptions which differs from the Power Architecture Book E definition. When a TEA (either a DTLB error or DSI, or a Machine Check (if MSR[EE]=0)) occurs in conjunction with an enabled DAC or linked DAC/IAC on a load or store class instruction, the Debug Interrupt takes priority, and the saved PC value points to the load or store class instruction, rather than to the next instruction. In addition, the MMU MAS registers are updated due to the DTLB event.

⁴ Unused on cacheless cores

5.9 Interrupt Processing

When an interrupt is taken, the processor uses SRR0/SRR1 for non-critical interrupts, CSRR0/CSRR1 for critical and machine check interrupts, and either CSRR0/CSRR1 or DSRR0/DSRR1 for debug interrupts to save the contents of the MSR and to assist in identifying where instruction execution should resume after the interrupt is handled.

When an interrupt occurs, one of SRR0/CSRR0/DSRR0 is set to the address of the instruction that caused the exception, or to the following instruction if appropriate.

SRR1 is used to save machine state (selected MSR bits) on non-critical interrupts and to restore those values when an **se_rfi** instruction is executed. CSRR1 is used to save machine status (selected MSR bits) on critical interrupts and to restore those values when an **se_rfci** instruction is executed. DSRR1 is used to save machine status (selected MSR bits) on debug interrupts when the Debug APU is enabled and to restore those values when an **se_rfdi** instruction is executed.

The Exception Syndrome register is loaded with information specific to the exception type. Some interrupt types can only be caused by a single exception type, and thus do not use an ESR setting to indicate the interrupt cause.

The Machine State register is updated to preclude unrecoverable interrupts from occurring during the initial portion of the interrupt handler. Specific settings are described in [Table 5-20](#).

For Alignment or Data Storage interrupts, the Data Exception Address Register (DEAR) is loaded with the address which caused the interrupt to occur.

For Machine Check interrupts, the Machine Check Syndrome register is loaded with information specific to the exception type.

Instruction fetch and execution resumes, using the new MSR value, at a location specific to the exception type. The location is determined by the Interrupt Vector Prefix Register (IVPR), and an Interrupt Vector Offset Register (IVOR) value specific for each type of interrupt (see [Table 5-2](#)).

[Table 5-20](#) shows the MSR settings for different interrupt categories.

Table 5-20. MSR Setting Due to Interrupt

Bit(s)	MSR Definition	Reset Setting	Non-Critical Interrupt	Critical Interrupt	Debug Interrupt
5 (37)	UCLE	0	0	0	0
13 (45)	WE	0	0	0	0
14 (46)	CE	0	—	0	—/0 ¹
16 (48)	EE	0	0	0	—/0 ¹
17 (49)	PR	0	0	0	0
18 (50)	FP	0	0	0	0
19 (51)	ME	0	—	—	—
20 (52)	FE0	0	0	0	0
22 (54)	DE	0	—	—/0 ¹	0
23 (55)	FE1	0	0	0	0
26 (58)	IS	0	0	0	0
27 (59)	DS	0	0	0	0
30 (62)	RI	0	—	0	—/0 ²

Reserved and preserved bits are unimplemented and read as 0.

¹ Conditionally cleared based on control bits in HID0

² Cleared if the Debug APU is disabled, otherwise unaffected

5.9.1 Enabling and Disabling Exceptions

When a condition exists that may cause an exception to be generated, it must be determined whether the exception is enabled for that condition.

- System reset exceptions cannot be masked.
- A machine check exception can occur only if the machine check enable bit (MSR[ME]) is set, or if a non-maskable interrupt is received. If MSR[ME] is cleared, the processor goes directly into checkstop state when a machine check exception condition occurs, unless the machine check is the

result of a non-maskable interrupt. Individual machine check exceptions (other than non-maskable interrupts) can be enabled and disabled through bit(s) in the HID0 register.

- Asynchronous, maskable non-critical exceptions (such as the External Input) are enabled by setting MSR[EE]. When MSR[EE]=0, recognition of these exception conditions is delayed. MSR[EE] is cleared automatically when a non-critical or critical interrupt is taken to mask further recognition of conditions causing those exceptions.
- Asynchronous, maskable critical exceptions (such as Critical Input) are enabled by setting MSR[CE]. When MSR[CE]=0, recognition of these exception conditions is delayed. MSR[CE] is cleared automatically when a critical interrupt is taken to mask further recognition of conditions causing those exceptions. In addition, MSR[RI] is cleared to indicate that the CSRR0/1 registers contain information essential to exception recovery.
- Synchronous and asynchronous Debug exceptions are enabled by setting MSR[DE]. When MSR[DE]=0, recognition of these exception conditions is masked. MSR[DE] is cleared automatically when a Debug interrupt is taken to mask further recognition of conditions causing those exceptions. See [Chapter 8, “Debug Support,”](#) for more details on individual control of debug exceptions.

5.9.2 Returning from an Interrupt Handler

The return from interrupt (**se_rfi**), return from critical interrupt (**se_rfci**) and return from debug interrupt (**se_rfdi**) instructions perform context synchronization by allowing previously-issued instructions to complete before returning to the interrupted process. In general, execution of a return instruction ensures the following:

- All previous instructions have completed to a point where they can no longer cause an exception. This includes post-execute type exceptions.
- Previous instructions complete execution in the context (privilege and protection) under which they were issued.
- The **se_rfi** instruction copies SRR1 bits back into the MSR.
- The **se_rfci** instruction copies CSRR1 bits back into the MSR.
- The **se_rfdi** instruction copies DSRR1 bits back into the MSR.
- Instructions fetched after this instruction execute in the context established by this instruction.
- Program execution resumes at the instruction indicated by SRR0 for **se_rfi**, CSRR0 for **se_rfci** and DSRR0 for **se_rfdi**.

Note that the return instruction **se_rfi** may be subject to a Return type debug exception, and that the return from critical interrupt instruction **se_rfci** may be subject to a Critical Return type debug exception. For a complete description of context synchronization, refer to Power Architecture Book E Specification.

5.10 Process Switching

The following instructions are useful for restoring proper context during process switching:

- The **msync** instruction orders the effects of data memory instruction execution. All instructions previously initiated appear to have completed before the **msync** instruction completes, and no subsequent instructions appear to be initiated until the **msync** instruction completes.
- The **isync** instruction waits for all previous instructions to complete and then discards any fetched instructions, causing subsequent instructions to be fetched (or refetched) from memory and to execute in the context (privilege, translation, and protection) established by the previous instructions.
- The **stwcx.** instructions clears any outstanding reservations, ensuring that a load and reserve instruction in an old process is not paired with a store conditional instruction in a new one.

Chapter 6

Core Complex Interfaces

This chapter describes the external interfaces to the e200 core complex. Signal descriptions as well as the data transfer protocols are documented in the following subsections.

The external interfaces encompass control and data signals supporting instruction and data transfers, support for interrupts, including vectored interrupt logic, reset support, power management interface signals, debug event signals, processor state information, Nexus/OnCE/JTAG interface signals, and a Test interface.

The memory portion of the e200 core interface is comprised of a pair of 32-bit wide system buses, one for instructions and the other for data in the e200z0h, and a unified bus on the e200z0. The data memory interface supports read and write transfers of 8, 16, 24, and 32 bits, supports misaligned transfers, and operates in a pipelined fashion. In the e200z0h the instruction memory interface supports read transfers of 16 and 32 bits, supports misaligned transfers, and operates in a pipelined fashion.

Single-beat and misaligned transfers are supported for read and write cycles. Incrementing burst transfers are supported for instruction prefetch operations.

Misaligned accesses are supported with one or more transfers to a bus interface. If an access is misaligned, but is contained within an aligned 32-bit word, the core performs a single transfer, and the memory interface is responsible for delivering (reads) or accepting (writes) the data corresponding to the size and byte enable signals aligned according to the low order two address bits. If an access is misaligned and crosses a 32-bit boundary, the bus interface unite (BIU) performs a pair of transfers beginning at the effective address for the first transfer, along with appropriate byte enables, and for the second transfer the address is incremented to the next 32-bit boundary, and the size and byte enable signals are driven to correspond to the number of remaining bytes to be transferred.

6.1 Signal Index

This section contains an index of the e200 signals.

The following prefixes are used for e200 signal mnemonics:

- *m* denotes master clock and reset signals
- *p* denotes processor or core-related signals
- *j* denotes JTAG mode signals
- *jd* denotes JTAG and Debug mode signals
- *ipt* denotes Scan and Test Mode signals
- *nex* denotes Nexus2 signals. Nexus signals support an optional Nexus2 and Nexus3 block on the e200z1 core.

NOTE

The “_b” suffix denotes an active low signal. Signals without the active-low suffix are active high.

[Figure 6-1](#) and [Figure 6-2](#) group core bus and control signals by function.




Figure 6-1. e200z0h Signal Groups




Figure 6-2. e200z0 Signal Groups

Table 6-1 shows e200 external signal function and type, signal definition, and reset value. Signals are presented in functional groups.

Table 6-1. External Interface Signal Definitions

Signal Name	Type	Reset Value	Definition
Clock and Reset-Related Signals			
<i>m_clk</i>	I	—	Global system clock
<i>m_por</i>	I	—	Power-on reset
<i>p_reset_b</i>	I	—	Processor reset input
<i>p_resetout_b</i>	O	—	Processor reset output
<i>p_rstbase</i> [0:29]	I	—	Reset exception handler base address
Memory Interface Signals			
<i>p_d_hmaster</i> [3:0], <i>p_i_hmaster</i> [3:0]	O	—	Master ID
<i>p_d_haddr</i> [31:0], <i>p_i_haddr</i> [31:0]	O	—	Address buses
<i>p_d_hwrite</i> , <i>p_i_hwrite</i> *	O	0	Write signal (always driven low for <i>p_i_hwrite</i>)
<i>p_d_hprot</i> [5:0], <i>p_i_hprot</i> [5:0]	O	—	Protection Codes
<i>p_d_htrans</i> [1:0], <i>p_i_htrans</i> [1:0]	O	—	Transfer Type
<i>p_d_hburst</i> [2:0], <i>p_i_hburst</i> [2:0]	O	—	Burst Type
<i>p_d_hsize</i> [1:0], <i>p_i_hsize</i> [1:0]	O	—	Transfer Size
<i>p_d_hunalign</i> , <i>p_i_hunalign</i>	O	—	Indicates the current access is a misaligned access.
<i>p_d_hbstrb</i> [3:0], <i>p_i_hbstrb</i> [3:0]	O	0	Byte strobes
<i>p_d_hrdata</i> [31:0], <i>p_i_hrdata</i> [31:0]	I	—	Read data buses
<i>p_d_hwdata</i> [31:0]	O	—	Write data bus
<i>p_d_hready</i> , <i>p_i_hready</i>	I	—	Transfer Ready
<i>p_d_hresp</i> [2:0], <i>p_i_hresp</i> [1:0]	I	—	Transfer Response
Interrupt Interface Signals			
<i>p_nmi_b</i>	I	—	Non-maskable interrupt request
<i>p_extint_b</i>	I	—	External Input interrupt request
<i>p_critint_b</i>	I	—	Critical Input interrupt request
<i>p_avec_b</i>	I	—	Autovector request Use internal interrupt vector offset
<i>p_voffset</i> [0:9]	I	—	Interrupt vector offset for vectored interrupts
<i>p_iack</i>	O	0	Interrupt Acknowledge. Indicates an interrupt is being acknowledge.
<i>p_ipend</i>	O	0	Interrupt Pending. Indicates an interrupt is pending internally.
<i>p_mcp_b</i>	I	—	Machine Check input request

Table 6-1. External Interface Signal Definitions (continued)

Signal Name	Type	Reset Value	Definition
Misc. CPU Signals			
<i>p_pid0</i> [0:7]	O	0	PID0[24:31] outputs
<i>p_hid1_sysctl</i> [0:7]	O	0	HID1[16:23] outputs
CPU Reservation Signals			
<i>p_rsrv</i>	O	0	Reservation status
<i>p_rsrv_clr</i>	I	—	Clear Reservation flag
CPU State Signals			
<i>p_pstat</i> [0:5]	O	0	Indicates processor status
<i>p_EE</i> , <i>p_DE</i> , <i>p_CE</i> , <i>p_ME</i>	O	0	Reflect the values of these MSR bits
<i>p_mcp_out</i>	O	0	Indicates a machine check has occurred
<i>p_chkstop</i>	O	0	Indicates a checkstop has occurred
<i>p_doze</i>	O	0	Indicates low-power doze mode of operation
<i>p_nap</i>	O	0	Indicates low-power nap mode of operation
<i>p_sleep</i>	O	0	Indicates low-power sleep mode of operation
<i>p_wakeup</i>	O	0	Indicates to external clock control module to enable clocks and exit from low-power mode
<i>p_halt</i>	I	—	CPU halt request
<i>p_halted</i>	O	0	CPU halted
<i>p_stop</i>	I	—	CPU stop request
<i>p_stopped</i>	O	0	CPU stopped
<i>p_waiting</i>	O	0	CPU waiting
CPU Debug Event Signals			
<i>p_ude</i>	I	—	Unconditional Debug Event
<i>p_devt1</i>	I	—	Debug Event 1 input
<i>p_devt2</i>	I	—	Debug Event 2 input
Debug/Emulation Support Signals (Nexus 1/OnCE)			
<i>jd_en_once</i>	I	—	Enable full OnCE operation
<i>jd_debug_b</i>	O	1	Indicates processor has entered debug session
<i>jd_de_b</i>	I	—	Debug request
<i>jd_de_en</i>	O	0	Active -high output enable for DE_b open-drain IO cell
<i>jd_mclk_on</i>	I	—	Indicates the system clock controller is actively toggling <i>m_clk</i>
<i>jd_watchpt</i> [0:5]	O	0	Indicate an address watchpoint has occurred
Development Support Signals (Nexus 2)			

Table 6-1. External Interface Signal Definitions (continued)

Signal Name	Type	Reset Value	Definition
<i>nex_mcko</i>	O	—	Nexus 2/3 Clock Output
<i>nex_rdy_b</i>	O	—	Nexus 2/3 Ready Output
<i>nex_evto_b</i>	O	—	Nexus 2/3 Event-Out Output
<i>nex_evti_b</i>	I	—	Nexus 2/3 Event-In Input
<i>nex_mdo</i> [n:0]	O	—	Nexus 2/3 Message Data Output
<i>nex_mseo_b</i> [1:0]	O	—	Nexus 2/3 Message Start/End Output
JTAG-Related Signals			
<i>j_trst_b</i>	I	—	JTAG test reset from pad
<i>j_tclk</i>	I	—	JTAG test clock from pad
<i>j_tms</i>	I	—	JTAG test mode select from pad
<i>j_tdi</i>	I	—	JTAG test data input from pad
<i>j_tdo</i>	O	0	JTAG test data out to master controller or pad
<i>j_tdo_en</i>	O	0	Enables TDO output buffer
<i>j_tst_log_rst</i>	O	0	Indicates Test-Logic-Reset state of JTAG controller
<i>j_capture_ir</i>	O	0	Indicates Capture_IR state of JTAG controller
<i>j_update_ir</i>	O	0	Indicates Update_IR state of JTAG controller
<i>j_shift_ir</i>	O	0	Indicates Shift_IR state of JTAG controller
<i>j_capture_dr</i>	O	0	Indicates parallel test data register load state of JTAG controller
<i>j_shift_dr</i>	O	0	Indicates the TAP controller is in shift DR state
<i>j_update_gp_reg</i>	O	0	Updates JTAG controller test data register
<i>j_rti</i>	O	0	JTAG controller run-test-idle state
<i>j_key_in</i>	I	—	Input for providing data to be shifted out during Shift_IR state when <i>jd_en_once</i> is negated
<i>j_en_once_regsel</i>	O	0	external Enable Once register select
<i>j_nexus_regsel</i>	O	0	external Nexus register select
<i>j_sncr_regsel</i>	O	0	external Shared Nexus Control register select
<i>j_lsrl_regsel</i>	O	0	external LSRL register select

Table 6-1. External Interface Signal Definitions (continued)

Signal Name	Type	Reset Value	Definition
<i>j_gp_regsel</i> [0:11]	O	0	General-purpose external JTAG register select
<i>j_id_sequence</i> [0:1]	I	—	JTAG ID Register (2 MSBs of sequence field)
<i>j_id_version</i> [0:3]	I	—	JTAG ID Register Version Field
<i>j_serial_data</i>	I	—	Serial data from external JTAG registers
Test Primary Input/Output Signals			
Test Control Interface ¹	—	—	Test Mode determination
Scan Test Interface ¹	—	—	Scan Configuration and Testing
Memory BIST Interface ¹	—	—	Memory BIST Configuration and Testing

¹ Please refer to the e200 Test Guide for information on the Test signals

6.2 Internal Interface Signals

Table 6-1 shows e200 internal signal function and type, signal definition, and reset value. Signals are presented in functional groups. Note that these signals are for reference purposes and their functionality is beyond the scope of this document.

Table 6-2. Internal Interface Signal Definitions

Signal Name	Type	Reset Value	Definition
Data Memory Interface Signals			
<i>p_d_addr</i> [0:31]	O	—	Address bus
<i>p_d_rw_b</i>	O	1	Read/write
<i>p_d_tc</i> [0:1]	O	—	Transfer Code
<i>p_d_ttype</i> [0:3]	O	—	Transfer Type
<i>p_d_tsiz</i> [0:2]	O	—	Transfer Size
<i>p_d_seq_b</i>	O	1	Indicates the current access is in sequential address order from the last access.
<i>p_d_misal_b</i>	O	1	Indicates the current data access is the first portion of a misaligned access.
<i>p_d_err_kill</i>	O	1	Indicates the current access will cause an abort if terminated with error.
<i>p_d_treq_b</i>	O	1	Transfer Request Indicates a request for a bus cycle.
<i>p_d_tbusy_b</i>	O	1	Transfer Busy Indicates a bus cycle is in progress.
<i>p_d_abort_b</i>	O	1	Aborts a requested access.

Table 6-2. Internal Interface Signal Definitions (continued)

Signal Name	Type	Reset Value	Definition
<i>p_d_data_in</i> [0:31]	I	—	Input data bus
<i>p_d_data_out</i> [0:31]	O	—	Output data bus
<i>p_d_ta_b</i>	I	—	Transfer Acknowledge
<i>p_d_te_b</i>	I	—	Transfer Error
<i>p_d_bus_werr</i>	I	—	Buffered Write Bus Error
<i>p_d_tmiss_b</i>	I	—	Translation Miss
<i>p_d_boerr_b</i>	I	—	Byte Ordering Error
<i>p_d_xte_b</i>	I	—	Precise External Termination Error
<i>p_d_xfail_b</i>	I	—	Store Exclusive Failure
Instruction Memory Interface Signals			
<i>p_i_addr</i> [0:31]	O	—	Address bus
<i>p_i_tc</i> [0:4]	O	—	Transfer Code
<i>p_i_tsiz</i> [0:2]	O	—	Transfer Size
<i>p_i_seq_b</i>	O	1	Indicates the current access is in sequential address order from the last access. For sequential instruction fetches.
<i>p_i_err_kill</i>	O	1	Indicates the current access will cause an abort if terminated with error.
<i>p_i_treq_b</i>	O	1	Transfer Request Indicates a request for a bus cycle.
<i>p_i_tbusy_b</i>	O	1	Transfer Busy Indicates a bus cycle is in progress.
<i>p_i_abort_b</i>	O	1	Aborts a requested access.
<i>p_i_data_in</i> [0:31]	I	—	Input data bus
<i>p_i_ta_b</i>	I	—	Transfer Acknowledge
<i>p_i_te_b</i>	I	—	Transfer Error
<i>p_i_xte_b</i>	I	—	Precise External Termination Error
<i>p_ifsiz</i>	I	—	Instruction Fetch Size Indicates the port size of the device being accessed.
SPR Interface Signals			
<i>p_sprnum</i> [0:9]	O	—	Global SPR address bus
<i>p_spr_out</i> [0:31]	O	—	Global SPR write bus
<i>p_spr_in</i> [0:31]	I	—	Global SPR read bus
<i>p_rd_spr</i>	O	0	SPR read control
<i>p_wr_spr</i>	O	0	SPR write control

Table 6-2. Internal Interface Signal Definitions (continued)

Signal Name	Type	Reset Value	Definition
Misc. CPU Signals			
<i>p_pid0[0:7]</i>	O	0	PID0[24:31] outputs
<i>p_pid0_updt</i>	O	0	PID0 update status
<i>p_d_cmbusy, p_i_cmbusy</i>	I	—	BIU busy
Test Primary Input/Output Signals			
Test Control Interface	—	—	Test Mode determination
Scan Test Interface	—	—	Scan Configuration and Testing
Memory BIST Interface	—	—	Memory BIST Configuration and Testing

6.3 Signal Descriptions

The following paragraphs provide descriptions of the external signals.

6.3.1 e200 Processor Clock (*m_clk*)

The *m_clk* input is the synchronous clock source for the e200 processor core.

Because e200 is designed for static operation, *m_clk* can be gated off to lower power dissipation (for example, during low-power stopped states).

6.3.2 Reset-Related Signals

e200 supports several reset input signals for the CPU and JTAG/OnCE control logic: *m_por*, *p_reset_b*, *p_resetout_b* and *j_trst_b*. The reset domains have been partitioned such that the CPU *p_reset_b* signal does not affect JTAG/OnCE logic and *j_trst_b* does not affect processor logic. It is possible and desirable to access OnCE registers while the processor is running or in reset. Alternatively, it is also possible and desirable to assert *j_trst_b* and clear the JTAG/OnCE logic without affecting the state of the processor.

The synchronization logic between the processor and debug module requires an assertion of either *j_trst_b* or *m_por* during initial processor power-up reset in order to ensure proper operation. If the pin associated with the *j_trst_b* input is designed with a pull-up resistor and left floating, then assertion of *m_por* is required during the initial power-on processor reset. Similarly, for those systems which do not have a power-on reset circuit and choose to tie *m_por* low, it is required to assert *j_trst_b* during processor power-up reset. Once a power-up reset has been achieved, the two resets can be asserted independently.

A reset output signal *p_resetout_b* is also provided.

A set of input signals (*p_rstbase[0:29]*) are provided to relocate the reset exception handler to allow for flexible placement of boot code.

These signals are described in detail in the following sub-sections.

6.3.2.1 Power-On Reset (*m_por*)

The *m_por* signal is the power-on reset input for the e200 processor. This signal serves the following purposes:

- *m_por* is “ORed” with the *j_trst_b* function and the resulting signal clears the JTAG TAP controller and associated registers as well as the OnCE state machine. This is an asynchronous clear with a short assertion time requirement.
- *m_por* is “ORed” with the *p_reset_b* function and the resulting signal clears certain CPU registers. This is an asynchronous clear with a short assertion time requirement.

6.3.2.2 Reset (*p_reset_b*)

The *p_reset_b* input is the active-low reset input for the e200 processor. *p_reset_b* is treated as an asynchronous input and is sampled by the clock control logic in the e200 debug module.

6.3.2.3 Reset Out (*p_resetout_b*)

The *p_resetout_b* output is an active-low reset output control signal from the e200 core. *p_resetout_b* is conditionally asserted by Debug control logic (Section 8.3.2.1, “Debug Control Register 0 (DBCRO0)”). *p_resetout_b* is not asserted by *p_reset_b*.

6.3.2.4 Reset Base (*p_rstbase*[0:29])

The *p_rstbase*[0:29] inputs are provided to allow system integrators to be able to specify/relocate the base address of the reset exception handler. These inputs are used to form the upper 30 bits of the instruction access following negation of reset which is used to fetch the initial instruction of the reset exception handler. These bits should be driven to a value corresponding to the desired boot memory device in the system. These inputs must remain stable in a window beginning two clocks prior to the negation of reset and extending into the cycle in which the reset vector fetch is initiated.

The initial instruction fetch occurs to the location [*p_rstbase*[0:29]] || 2'b00.

6.3.2.5 JTAG/OnCE Reset (*j_trst_b*)

The *j_trst_b* signal (referred to in the *IEEE 1149.1 JTAG Specification* as the TRST* signal) is an asynchronous reset with a short assertion time requirement. It is “ORed” with the *m_por* function and the resulting signal clears the OnCE TAP controller and associated registers as well as the OnCE state machine.

6.3.3 Address and Data Buses

Dual instruction and data interfaces are provided by the CPU. They are described together, with appropriate differences denoted.

6.3.3.1 Address Bus ($p_d_haddr[31:0]$, $p_i_haddr[31:0]$)

These outputs provide the address for a bus transfer. Per the AHB definition, $p_d[i]_haddr[31]$ is the MSB and $p_d[i]_haddr[0]$ is the LSB.

6.3.3.2 Read Data Bus ($p_d_hrdata[31:0]$, $p_i_hrdata[31:0]$)

These inputs provide data to the CPU on read transfers. The data read data bus can transfer 8, 16, 24, or 32 bits of data per bus transfer. The instruction read data bus can transfer 16 or 32 bits of data per bus transfer. Instruction transfers do not use the 8-bit and 24-bit capability. Per AHB definition, $p_d[i]_hrdata[31]$ is the MSB and $p_hrdata[0]$ is the LSB. [Table 6-3](#) shows the relationship of byte addresses to read data bus signals.

Table 6-3. $p_hrdata[31:0]$ Byte Address Mappings

Memory Byte Address	Wired to p_d_hrdata Bits
00	7:0
01	15:8
10	23:16
11	31:24

6.3.3.3 Write Data Bus ($p_d_hwdata[31:0]$)

These outputs transfer data from the CPU on write transfers. The write data bus can transfer 8, 16, 24, or 32 bits of data per bus transfer. Per AHB definition, $p_d_hwdata[31]$ is the MSB and $p_d_hwdata[0]$ is the LSB. [Figure 6-4](#) shows the relationship of byte addresses to write data bus signals.

Table 6-4. $p_d_hwdata[31:0]$ Byte Address Mappings

Memory Byte Address	Wired to p_d_hwdata Bits
00	7:0
01	15:8
10	23:16
11	31:24

6.3.4 Transfer Attribute Signals

The following paragraphs describe the transfer attribute signals, which provide additional information about the bus transfer cycle. Transfer attributes are driven with address at the beginning of a bus transfer.

6.3.4.1 Transfer Type ($p_d_htrans[1:0]$, $p_i_htrans[1:0]$)

The processor drives these signals to indicate the current transfer type. Table 6-5 shows $p_d,i_htrans[1:0]$ encoding.

Table 6-5. $p_d,i_htrans[1:0]$ Transfer Type Encoding

$p_d,i_htrans[1]$	$p_d,i_htrans[0]$	Access type
0	0	IDLE—no data transfer is required
0	1	BUSY—Master is busy, burst transfer continues. (encoding not used by e200)
1	0	NONSEQ—indicates the first transfer of a burst, or a single transfer. Address and control signals are unrelated to the previous transfer
1	1	SEQ—indicates the continuation of a burst. Address and control signals are related to the previous transfer. Control signals are the same, Address has been incremented by the size of the data transferred (optionally wrapped)

If the $p_d,i_htrans[1:0]$ encoding is not IDLE or BUSY, a transfer is being requested. e200 does not utilize the BUSY encoding, and does not present this type of transfer to a bus slave. Slaves must terminate IDLE transfers with a zero wait-state OKAY response and ignore the (non-existent) transfer.

6.3.4.2 Write (p_d_hwrite , p_i_hwrite)

This output signal defines the data transfer direction for the current bus cycle. A high (logic one) level indicates a write cycle, and a low (logic zero) level indicates a read cycle. For p_i_hwrite , the signal is internally driven low for all IAHB transfers.

6.3.4.3 Transfer Size ($p_d_hsize[1:0]$, $p_i_hsize[1:0]$)

The $p_d,i_hsize[1:0]$ signals indicate the data size for a bus transfer. Table 6-6 shows the definitions of the $p_d,i_hsize[1:0]$ encodings. For misaligned transfers, the transfer size may indicate a size larger than the requested size to ensure that all asserted byte strobes are contained within the “container” defined by $p_d,i_hsize[1:0]$. Refer to Table 6-10 and Table 6-11 for $p_d,i_hsize[1:0]$ encodings used for aligned and misaligned transfers.

Table 6-6. $p_d,i_hsize[1:0]$ Transfer Size Encoding

$p_d,i_hsize[1:0]$	Transfer Size
00	Byte
01	Halfword (2 bytes)
10	Word (4 bytes)
11	Reserved

6.3.4.4 Burst Type ($p_d_hburst[2:0]$, $p_i_hburst[2:0]$)

The $p_d[i]_hburst[2:0]$ signals indicate the burst type for a bus transfer. Table 6-7 shows the definitions of the $p_d[i]_hburst[2:0]$ encodings.

Table 6-7. $p_d[i]_hburst[2:0]$ Burst Type Encoding

$p_hburst[2:0]$	Burst Type
000	SINGLE—No burst, single beat only
001	INCR—Incrementing burst of unspecified length
010	WRAP4—4-beat wrapping burst—Unused
011	INCR4—4-beat incrementing burst—Unused
100	WRAP8—8-beat wrapping burst—Unused
101	INCR8—8-beat incrementing burst—Unused
110	WRAP16—16-beat wrapping burst—Unused
111	INCR16—16-beat incrementing burst—Unused

e200 only utilizes SINGLE and INCR (for instruction) burst types. In addition, all INCR bursts are of word size aligned to word boundaries.

6.3.4.5 Protection Control ($p_d_hprot[5:0]$, $p_i_hprot[5:0]$)

e200 drives the $p_d[i]_hprot[5:0]$ signals to indicate the type of access for the current bus cycle. $p_d[i]_hprot[0]$ indicates instruction/data, $p_d[i]_hprot[1]$ indicates user/supervisor. $p_d[i]_hprot[5]$ indicates whether the access is Exclusive (such as for a **lwarx** or **stwcx.**). $p_d[i]_hprot[4:2]$ (Allocate, Cacheable, Bufferable) are used to indicate particular cache attributes for the access and are driven to default values of 3'b000. Table 6-8 shows the definitions of the $p_d[i]_hprot[5:0]$ signals.

Table 6-8. $p_d[i]_hprot[5:0]$ Protection Control Encoding

$p_hprot[5]$	$p_hprot[4]$	$p_hprot[3]$	$p_hprot[2]$	$p_hprot[1]$	$p_hprot[0]$	Transfer Type
—	—	—	—	—	0	Instruction Access
—	—	—	—	—	1	Data Access
—	—	—	—	0	—	User mode access
—	—	—	—	1	—	Supervisor mode access
0	—	—	—	—	—	Not Exclusive
1	—	—	—	—	—	Exclusive Access

Note that all signals are provided on both I and D ports, although they do not all change state (for example, *p_d_hprot0* is always high, etc.).

6.3.5 Byte Lane Specification

Read transactions transfer from 1 to 4 bytes of data on the *p_[d,i]_hrdata[31:0]* bus. The byte lanes involved in the transfer are determined by the starting byte number specified by the lower address bits in conjunction with the transfer size and byte strobes. Addressing of the byte lanes is shown big-endian (left to right). The byte of memory corresponding to address 0 is connected to B0 (*p_[d,i]_h{r,w}data[7:0]*) and the byte of memory corresponding to address 3 is connected to B3 (*p_[d,i]_h{r,w}data[31:24]*). The CPU internally permutes read data as required for the current access. Misaligned transfers are indicated with the *p_[d,i]_hunalign* signal to indicate that byte strobes do not correspond exactly to size and low-order address bits.

6.3.5.1 Unaligned Access (*p_d_hunalign*, *p_i_hunalign*)

The *p_[d,i]_hunalign* output signal indicates that the current access is a misaligned access. This signal is asserted for misaligned data accesses. The timing of this signal is approximately the same as address timing. When *p_[d,i]_hunalign* is asserted, the *p_[d,i]_hbstrb[3:0]* byte strobe signals indicate the selected bytes involved in the current portion of the misaligned access, which may not include all bytes defined by the size and low-order address signals. Aligned transfers also assert the byte strobes, but in a manner corresponding to size and low order address bits.

6.3.5.2 Byte Strobes (*p_d_hbstrb[3:0]*, *p_i_hbstrb[3:0]*)

The *p_[d,i]_hbstrb[3:0]* byte strobe signals indicate the selected bytes involved in the current transfer. For a misaligned access, the current transfer may not include all bytes defined by the size and low-order address signals. For aligned transfers, the byte strobe signals correspond to the bytes defined by the size and low-order address signals. [Table 6-9](#) shows the relationship of byte addresses to the byte strobe signals.

Table 6-9. *p_hbstrb[3:0]* to Byte Address Mappings

Memory Byte Address	Wired to <i>p_h{r,w}data</i> Bits	Corresponding Byte Strobe Signal
00	7:0	<i>p_[d,i]_hbstrb[0]</i>
01	15:8	<i>p_[d,i]_hbstrb[1]</i>
10	23:16	<i>p_[d,i]_hbstrb[2]</i>
11	31:24	<i>p_[d,i]_hbstrb[3]</i>

Table 6-10 lists all of the data transfer permutations. Note that misaligned data requests which cross a 32-bit boundary are broken up into two separate bus transactions, and the address value and the size encoding for the first transfer is not modified. The table is arranged in a big-endian fashion. e200 performs the proper byte routing internally .

Table 6-10. Byte Strobe Assertion for Transfers

Program Size and Byte Offset	A(1:0)	HSIZE [1:0]	Data Bus Byte Strobes				HUNALIGN
			B0	B1	B2	B3	
Byte @00	0 0	0 0	X	—	—	—	0
Byte @01	0 1	0 0	—	X	—	—	0
Byte @10	1 0	0 0	—	—	X	—	0
Byte @11	1 1	0 0	—	—	—	X	0
Half @00	0 0	0 1	X	X	—	—	0
Half @01	0 1	1 0 [#]	—	X	X	—	1
Half @10	1 0	0 1	—	—	X	X	0
Half @11 (2 bus transfers)	1 1	0 1 [*]	—	—	—	X	1
	0 0	0 0	X	—	—	—	0
Word @00	0 0	1 0	X	X	X	X	0
Word @01 (2 bus transfers)	0 1	1 0 [*]	—	X	X	X	1
	0 0	0 0	X	—	—	—	0
Word @10 (2 bus transfers)	1 0	1 0 [*]	—	—	X	X	1
	0 0	0 1	X	X	—	—	0
Word @11 (2 bus transfers)	1 1	1 0 [*]	—	—	—	X	1
	0 0	1 0 [#]	X	X	X	—	1

Note:

“X” indicates byte lanes involved in the transfer; Other lanes contain driven but unused data.

[#] These misaligned transfers drive size according to the size of the power of two aligned “container” in which the byte strobes are asserted.

^{*} These misaligned cases drive request size according to the size specified by the load or store instruction.

Table 6-11 shows the final layout in memory for data transferred from a 32-bit GPR containing the bytes ‘E F G H’ to memory. Misaligned accesses which cross a word boundary are broken into a pair of accesses by the CPU.

Table 6-11. Big-Endian Memory Storage

Program Size and Byte Offset	A(2:0)	HSIZE (1:0)	Even Word—0				Odd Word—1			
			B0	B1	B2	B3	B0	B1	B2	B3
Byte @000	0 0 0	0 0	H	—	—	—	—	—	—	—
Byte @001	0 0 1	0 0	—	H	—	—	—	—	—	—
Byte @010	0 1 0	0 0	—	—	H	—	—	—	—	—
Byte @011	0 1 1	0 0	—	—	—	H	—	—	—	—
Byte @100	1 0 0	0 0	—	—	—	—	H	—	—	—
Byte @101	1 0 1	0 0	—	—	—	—	—	H	—	—
Byte @110	1 1 0	0 0	—	—	—	—	—	—	H	—
Byte @111	1 1 1	0 0	—	—	—	—	—	—	—	H
B. E. Half @000	0 0 0	0 1	G	H	—	—	—	—	—	—
B. E. Half @001	0 0 1	1 0 [#]	—	G	H	—	—	—	—	—
B. E. Half @010	0 1 0	0 1	—	—	G	H	—	—	—	—
B. E. Half @011	0 1 1 1 0 0	0 1 0 0	—	—	—	G	H	—	—	—
B. E. Half @100	1 0 0	0 1	—	—	—	—	G	H	—	—
B. E. Half @101	1 0 1	1 0 [#]	—	—	—	—	—	G	H	—
B. E. Half @110	1 1 0	0 1	—	—	—	—	—	—	G	H
B.E. Half @111	1 1 1	0 1	—	—	—	—	—	—	—	G
	+ 0 0 0 (next word)	0 0	H	—	—	—	—	—	—	—
B. E. Word @000	0 0 0	1 0	E	F	G	H	—	—	—	—
B. E. Word @001	0 0 1 1 0 0	1 0 [*] 0 0	—	E	F	G	H	—	—	—
B. E. Word @010	0 1 0 1 0 0	1 0 [*] 0 1	—	—	E	F	G	H	—	—
B. E. Word @011	0 1 1 1 0 0	1 0 [*] 1 0 [#]	—	—	—	E	F	G	H	—
B. E. Word @100	1 0 0	1 0	—	—	—	—	E	F	G	H
B. E. Word @101	1 0 1	1 0 [*]	—	—	—	—	—	E	F	G
	+ 0 0 0 (next word)	0 0	H	—	—	—	—	—	—	—
B. E. Word @110	1 1 0	1 0 [*]	—	—	—	—	—	—	E	F
	+ 0 0 0 (next word)	0 1	G	H	—	—	—	—	—	—

Table 6-11. Big-Endian Memory Storage (continued)

Program Size and Byte Offset	A(2:0)	HSIZE (1:0)	Even Word—0				Odd Word—1			
			B0	B1	B2	B3	B0	B1	B2	B3
B. E. Word @111	1 1 1	1 0*	—	—	—	—	—	—	—	E
	+ 0 0 0 (next word)	1 0#	F	G	H	—	—	—	—	—

Notes:

Assumes a 32-bit GPR contains 'E F G H'

These misaligned transfers drive size according to the size of the power of two aligned “container” in which the byte strobes are asserted.

* These misaligned cases drive request size according to the size specified by the load or store instruction.

6.3.6 Transfer Control Signals

The following paragraphs describe the transfer control signals.

6.3.6.1 Transfer Ready (*p_d_hready*, *p_i_hready*)

The *p_d_hready* input signal indicates completion of a requested transfer operation. An external device asserts *p_d_hready* to terminate the transfer. The *p_d_hresp*[2:0] signals indicate status of the transfer.

6.3.6.2 Transfer Response (*p_d_hresp*[2:0], *p_i_hresp*[1:0])

The *p_d_hresp*[2:0] and *p_i_hresp*[1:0] signals indicate status of a terminating transfer on the respective interfaces. Table 6-12 shows the definitions of the *p_d_hresp*[2:0] and *p_i_hresp*[1:0] encodings.

Table 6-12. *p_d_hresp*[2:0] Transfer Response Encoding

<i>p_d_hresp</i> [2:0]	Response Type
000	OKAY—transfer terminated normally
001	ERROR—transfer terminated abnormally
010	Reserved (RETRY not supported in AHB-Lite protocol)
011	Reserved (SPLIT not supported in AHB-Lite protocol)
100	XFAIL—Exclusive store failed (stwcx. did not completed successfully)
101	Reserved
110	Reserved
111	Reserved

Table 6-13. *p_i_hresp*[1:0] Transfer Response Encoding

<i>p_i_hresp</i> [1:0]	Response Type
00	OKAY—transfer terminated normally
01	ERROR—transfer terminated abnormally
10	Reserved (RETRY not supported in AHB-Lite protocol)
11	Reserved (SPLIT not supported in AHB-Lite protocol)

The ERROR and XFAIL responses are required to be two cycle responses. In this case, the ERROR or XFAIL responses must be signaled one cycle prior to assertion of *p_[d,i]_hready*, and must remain unchanged during the cycle *p_[d,i]_hready* is asserted.

The XFAIL response is signaled to the CPU.

6.3.7 Interrupt Signals

The following paragraphs describe the signals which control the interrupt functions. Interrupt request inputs *p_extint_b* and *p_critint_b* to the core are level sensitive, not edge-triggered, thus the interrupt controller module must keep the interrupt request as well as the *p_voffset* or *p_avec_b* inputs (as appropriate) asserted until the interrupt is serviced to guarantee that the CPU core recognizes the request. On the other hand, once a request is generated, there is no guarantee the CPU will not recognize the interrupt request even if the request is later removed. Interrupt requests must be held stable to avoid spurious responses. The interrupt inputs *p_nmi_b* and *p_mcp_b* are transition-sensitive and must be held asserted until acknowledged in order to be guaranteed to be recognized, although there is no guarantee the CPU will not recognize the interrupt request even if the request is later removed.

6.3.7.1 External Input Interrupt Request (*p_extint_b*)

This active-low signal provides the External Input interrupt request to the e200 core. *p_extint_b* is masked by the MSR[EE] bit. This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints relative to *m_clk* when the e200 core clock is running. This signal is level sensitive and must remain asserted to be guaranteed to be recognized.

6.3.7.2 Critical Input Interrupt Request (*p_critint_b*)

This active-low signal provides the Critical Input interrupt request to the e200 core. *p_critint_b* is masked by the MSR[CE] bit. This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints relative to *m_clk* when the e200 core clock is running. This signal is level sensitive and must remain asserted to be guaranteed to be recognized.

6.3.7.3 Non-Maskable Input Interrupt Request (*p_nmi_b*) on e200z0h

This active-low, transition sensitive signal provides a non-maskable interrupt request to the e200 core. This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints to *m_clk* when the e200 core clock is running. The *p_nmi_b* input is sampled on two consecutive *m_clk* periods to detect a transition from the negated to the asserted state. Initiation of exception processing for

the NMI is internally qualified with this transition, but must remain asserted low to be guaranteed to be recognized. Note that when the core is halted or stopped, without clocks, transitions on this signal are not immediately detected, but the *p_ipend* and *p_wakeup* signals are asserted to indicate to system logic that an interrupt is pending and so the clocks should be started, and the halt and stop inputs should be negated in order for the interrupt to be processed. Also, when the core is in the debug state, the internal *m_clk* is not running, so the *p_nmi_b* input does not guaranteed to be recognized until the core is released with a go+noexit or a go+exit OnCE command.

6.3.7.4 Interrupt Pending (*p_ipend*)

This active-high signal indicates that an asserted *p_extint_b*, *p_critint_b*, or *p_nmi_b* interrupt request input, or an enabled Timer facility interrupt (Watchdog, Fixed-Interval, or Decrementer) has been recognized internally by the core and is enabled by the appropriate bit in the MSR (*p_nmi_b* is never masked), and is asserted combinationally from the qualified interrupt request inputs. The *p_ipend* signal can be used to signal other bus masters or a bus arbiter that an interrupt condition is pending. External power management logic can use this output to control operation of the core and other logic or may use the *p_wakeup* signal similarly. Actual handling of the interrupt request may be delayed due to higher priority exceptions; assertion of *p_ipend* does not mean that exception processing for the interrupt has begun. The *p_nmi_b* input affects the *p_ipend* signal slightly differently; the *p_ipend* output asserts any time the *p_nmi_b* input is asserted.

6.3.7.5 Autovector (*p_avec_b*)

This active-low signal is asserted with either the *p_extint_b* or *p_critint_b* interrupt request to request use of the internal IVOR4 or IVOR0 values for obtaining an exception vector offset. If this signal is negated when a *p_extint_b* or *p_critint_b* interrupt is requested, an external vector offset is taken from the *p_voffset*[0:9] input signals. This signal is level sensitive and must remain asserted to be guaranteed to be recognized. This signal must be driven to a valid state during each clock cycle that either *p_extint_b* or *p_critint_b* is asserted.

6.3.7.6 Interrupt Vector Offset (*p_voffset*[0:9])

These input signals provide a vector offset to be used when exception processing begins for an incoming interrupt request. These signals are sampled along with the *p_extint_b* and *p_critint_b* interrupt request inputs, and must be driven to a valid value when either of these signals is asserted unless the *p_avec_b* signal is also asserted. If *p_avec_b* is asserted, these inputs are not used. The *p_voffset*[0:9] signals correspond to bits 20:29 of the exception handler address (the low order two bits 30:31 are forced to 00). The *p_voffset*[0:9] signals are level sensitive and must remain asserted to be guaranteed to be recognized correctly. In addition, these signals must be asserted concurrently with the *p_extint_b* and *p_critint_b* inputs when used.

6.3.7.7 Interrupt Vector Acknowledge (*p_iack*)

The *p_iack* output signal provide an interrupt vector acknowledge indicator to allow external interrupt controllers to be informed when a critical input or external input interrupt is being processed. The *p_iack*

signal is asserted after the cycle in which the *p_avec_b* and *p_voffset[0:9]* signals are sampled in preparation for exception processing. See [Table 6-13](#) and [Figure 6-32](#) for timing diagrams of operation.

6.3.7.8 Machine Check (*p_mcp_b*)

This active-low signal provides the Machine Check interrupt request to the e200 core. *p_mcp_b* is masked by the HID0[EMCP] bit. This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints to *m_clk* when the e200 core clock is running. The *p_mcp_b* input is sampled on two consecutive *m_clk* periods to detect a transition from the negated to the asserted state. It is internally qualified with this transition, but must remain asserted to be guaranteed to be recognized.

The *p_mcp_b* signal is not sampled while the e200 core is in the halted or stopped power management states, but is sampled while the CPU is in the waiting state. See [Section 6.3.11.3, “Processor Halted \(*p_halted*\),”](#) and [Section 6.3.11.5, “Processor Stopped \(*p_stopped*\).”](#) Also, when the core is in the debug state (as reflected on the *cpu_dbgack* internal state signal), the internal *m_clk* is not running, so the *p_mcp_b* input is not recognized until the core is released with a go+noexit or a go+exit OnCE command.

6.3.8 Processor Reservation Signals

The following sub-sections describe processor reservation signals associated with the **lwarx** and **stwcx..** instructions.

6.3.8.1 CPU Reservation Status (*p_rsrv*)

The active-high *p_rsrv* output signal is used to indicate that a reservation has been established by the execution of a **lwarx** instruction. This signal is set following the successful completion of a **lwarx**. This signal remains set until the reservation has been cleared. (Refer to [Section 3.4, “Memory Synchronization and Reservation Instructions”](#)). This signal is provided as a status indicator for specialized system applications only.

6.3.8.2 CPU Reservation Clear (*p_rsrv_clr*)

The active-high *p_rsrv_clr* input signal is used to clear a reservation that has been previously established. External reservation management logic may use this signal to implement reservation management policies which are outside of the scope of the CPU. (Refer to [Section 3.4, “Memory Synchronization and Reservation Instructions”](#)). This signal may be asserted independently of any bus transfer.

The *p_rsrv_clr* input signal is not intended for normal use in managing reservations. It is provided for specialized system applications. The normal bus protocol is used to manage reservations using external reservation logic in systems with multiple coherent bus masters, using the transfer type and transfer response signals. In single coherent master systems, no external logic is required, and the internal reservation flag is sufficient to support multi-tasking applications.

The *p_d_xfail_b* signal is provided to indicate success/failure of a **stwcx.** instruction as part of bus transfer termination using the XFAIL *p_d_hresp[2:0]* encoding. See [Section 7.2.3.7, “Store Exclusive Failure \(*p_d_xfail_b*\),”](#) for more detail on *p_d_xfail_b*.

6.3.9 Miscellaneous Processor Signals

The following paragraph describes several miscellaneous processor signals.

6.3.9.1 PID0 Outputs (*p_pid0*[0:7])

The active-high *p_pid0*[0:7] output signals are used to provide the current process ID in the Process ID Register 0 (PID0). These outputs correspond to the low order eight bits of PID0.

6.3.9.2 PID0 Update (*p_pid0_updt*)

The active-high *p_pid0_updt* signal is used to indicate that the Process ID Register 0 (PID0) is being updated by a *mtspr* instruction. This output asserts during the clock cycle the *p_pid0*[0:7] outputs are changing.

6.3.9.3 HID1 System Control (*p_hid1_sysctl*[0:7])

The active-high *p_hid1_sysctl*[0:7] output signals are used to provide a set of control output signals external to the CPU via values written to the HID1 special purpose register. These outputs change state following the rising edge of *m_clk*, and may need synchronization depending on actual use. See [Section 2.3.10, “Hardware Implementation Dependent Register 1 \(HID1\).”](#)

6.3.10 Processor State Signals

The following sub-sections describe processor internal state signals.

6.3.10.1 Processor Status (*p_pstat*[0:5])

These signals indicate the internal execution unit status. The timing is synchronous with the *m_clk*, so the indicated status may not apply to a current bus transfer. [Table 6-14](#) shows *p_pstat*[0:5] encoding.

Table 6-14. Processor Status Encoding¹

<i>p_pstat</i> [0:5]						Internal Processor Status
0	0	0	0	0	x	Execution Stalled
0	0	0	0	1	x	Execute Exception
0	0	0	1	0	x	Instruction Squashed
0	0	0	1	1	x	Reserved
0	0	1	0	0	x	Reserved
0	0	1	0	1	x	Reserved
0	0	1	1	0	x	Reserved
0	0	1	1	1	x	Processor in Waiting State
0	1	0	0	0	x	Processor in Halted state
0	1	0	0	1	x	Processor in Stopped state

Table 6-14. Processor Status Encoding¹ (continued)

<i>p_pstat</i> [0:5]						Internal Processor Status
0	1	0	1	0	x	Processor in Debug mode ²
0	1	0	1	1	x	Processor in Checkstop state
0	1	1	0	0	x	Reserved
0	1	1	0	1	x	Reserved
0	1	1	1	0	x	Reserved
0	1	1	1	1	x	Reserved
1	0	0	0	0	s	Complete Instruction ^{3,4}
1	0	0	0	1	0	Complete e_lmw , or e_stmw
1	0	0	1	0	1	Complete se_isync
1	0	0	1	1	0	Complete lwarx or stwcx .
1	0	1	0	0	0	Reserved
1	0	1	0	1	0	Reserved
1	0	1	1	0	0	Reserved
1	0	1	1	1	0	Reserved
1	1	0	0	0	0	Complete Branch Instruction e_bc , e_bcl , e_b , e_bl resolved as not taken
1	1	0	0	0	1	Complete Branch Instruction se_bc , se_bcl , se_b , se_bl resolved as not taken
1	1	0	0	1	0	Complete Branch Instruction e_bc , e_bcl , e_b , e_bl resolved as taken
1	1	0	0	1	1	Complete Branch Instruction se_bc , se_bcl , se_b , se_bl resolved as taken
1	1	0	1	1	1	Complete se_blr , se_brlr , se_bctr , se_bctrl (always taken)
1	1	1	0	0	0	Complete isel with condition false
1	1	1	0	1	0	Complete isel with condition true
1	1	1	1	0	x	Reserved
1	1	1	1	1	1	Complete se_rfi , se_rfci , or se_rfdi

¹ All encodings which do not appear in the table are reserved

² As reflected on the **cpu_dbgack** internal state signal

³ Except **rfi**, **rfci**, **rfdi**, **lmw**, **stmw**, **lwarx**, **stwcx.**, **isync**, **isel**, **se_rfi**, **se_rfci**, **se_rfdi**, **e_lmw**, **e_stmw**, **se_isel**, and Change of Flow Instructions

⁴ s—instruction size, 0=32-bit, 1=16-bit

6.3.10.2 Processor Exception Enable MSR Values (*p_EE*, *p_CE*, *p_DE*, *p_ME*)

These active-high output signals reflect the state of the corresponding MSR[EE,CE,DE,ME] bits. They may be used by external system logic to determine the set of enabled exceptions. These signals change state on execution of a **mtmsr**, **se_rfi**, **se_rfci**, **se_rfdi**, **wrtee**, or **wrteei** instruction, or during exception processing where one or more bits may be cleared during the exception processing sequence.

6.3.10.3 Processor Machine Check (*p_mcp_out*)

The active-high *p_mcp_out* output signal is asserted by the processor when a machine check condition has caused a syndrome bit to be set in the Machine Check Syndrome register. Refer to [Section 2.3.7, “Machine Check Syndrome Register \(MCSR\)”](#).

6.3.10.4 Processor Checkstop (*p_chkstop*)

The active-high *p_chkstop* output signal is asserted by the processor when a checkstop condition has occurred and the CPU has entered the checkstop state.

6.3.11 Power Management Control Signals

The following signals are provided for power management or other control functions by external control logic.

6.3.11.1 Processor Waiting (*p_waiting*)

The active-high *p_waiting* output signal is used to indicate that the processor has entered the Waiting state ([Section 8.1.2, “Waiting State”](#)).

6.3.11.2 Processor Halt Request (*p_halt*)

The active-high *p_halt* input signal is used to request the processor to enter the Halted state ([Section 8.1.3, “Halted State”](#)).

6.3.11.3 Processor Halted (*p_halted*)

The active-high *p_halted* output signal is used to indicate that the processor has entered the Halted state ([Section 8.1.3, “Halted State”](#)).

6.3.11.4 Processor Stop Request (*p_stop*)

The active-high *p_stop* input signal is used to request the processor to enter the Stopped state ([Section 8.1.4, “Stopped State”](#)).

6.3.11.5 Processor Stopped (*p_stopped*)

The active-high *p_stopped* output signal is used to indicate that the processor has entered the Stopped state ([Section 8.1.4, “Stopped State”](#)).

6.3.11.6 Low-Power Mode Signals (*p_doze*, *p_nap*, *p_sleep*)

The active-high *p_doze*, *p_nap*, and *p_sleep* output signals are asserted by the processor to reflect the settings of the HID0[DOZE], HID0[NAP], and HID0[SLEEP] control bits when the MSR[WE] bit is set.

These outputs may assert for one or more clock cycles. External logic can detect the asserted edge or level of these signals to determine which low-power mode has been requested and then place the e200 core and

peripherals in a low-power consumption state. The *p_wakeup* signal can be monitored to determine when to end the low-power condition.

The e200 core can be placed in a low-power state by forcing the *m_clk* input to a quiescent state, and brought out of low-power state by re-enabling *m_clk*.

6.3.11.7 Wakeup (*p_wakeup*)

The active-high *p_wakeup* output signal should be used by external logic to remove the e200 core and system logic from a low-power state. It also is used to indicate to the system clock controller that the *m_clk* input should be re-enabled for debug purposes. This signal is asynchronous to the system clock and should be synchronized to the system clock domain to avoid hazards.

p_wakeup asserts whenever the following occurs:

- A valid pending interrupt is detected by the core
- A request to enter debug mode is made by setting the DR bit in the OnCE control register (OCR) or via the assertion of the *jd_de_b* or *p_ude* input signals.
- The processor is in a debug session and the *jd_debug_b* output is asserted
- A request to enable the *m_clk* input has been made by setting the WKUP bit in the OnCE control register
- The *p_nmi_b* input is asserted

p_wakeup (or other system state) should be monitored to determine when to release the processor (and system if applicable) from a low-power state.

6.3.12 Debug Event Signals

The following interface signals are provided to signal debug events to the e200 core.

6.3.12.1 Unconditional Debug Event (*p_ude*)

The active-high *p_ude* input signal is used to request an unconditional debug event. This event is described in detail in [Section 8.2.12, “Unconditional Debug Event.”](#) This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints relative to *m_clk* when the e200 core clock is running. This signal is level sensitive and must be held asserted until acknowledged by software, or, when external debug mode is enabled, by assertion of the *jd_debug_b* output to be guaranteed to be recognized. In addition, only a transition from the negated state to the asserted state of the *p_ude* signal causes an event to occur. The level on this signal is used however to cause assertion of the *p_wakeup* output.

6.3.12.2 External Debug Event 1 (*p_dev1*)

The active-high *p_dev1* input signal is used to request an external debug event. This event is described in detail in [Section 8.2.11, “External Debug Event.”](#) This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints relative to *m_clk* when the e200 core clock is running. If the e200 core clock is disabled, this signal is not recognized. In addition, only a transition from

the negated state to the asserted state of the *p_devt1* signal causes an event to occur. It is intended to signal e200 related events that are generated while the CPU is active.

6.3.12.3 External Debug Event 2 (*p_devt2*)

The active-high *p_devt2* input signal is used to request an external debug event. This event is described in detail in [Section 8.2.11, “External Debug Event.”](#) This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints relative to *m_clk* when the e200 core clock is running. If the e200 core clock is disabled, this signal is not recognized. In addition, only a transition from the negated state to the asserted state of the *p_devt2* signal causes an event to occur. It is intended to signal e200 related events that are generated while the CPU is active.

6.3.13 Debug/Emulation (Nexus 1/OnCE) Support Signals

The following interface signals are provided to assist in implementing an On-Chip Emulation capability with a controller external to the e200 core.

Table 6-15. e200 Debug/Emulation Support Signals

Signal	Type	Description
<i>jd_en_once</i>	I	Enable full OnCE operation
<i>jd_debug_b</i>	O	Debug Session indicator
<i>jd_de_b</i>	I	Debug request
<i>jd_de_en</i>	O	DE_b active high output enable
<i>jd_mclk_on</i>	I	CPU clock is active indicator

6.3.13.1 OnCE Enable (*jd_en_once*)

The OnCE enable signal *jd_en_once* is used to enable the OnCE controller to allow certain instructions and operations to be executed. Assertion of this signal enables the full OnCE command set, as well as operation of control signals and OnCE Control register functions. When this signal is disabled, only the Bypass, ID and Enable_OnCE commands are executed by the Z5 OnCE unit, and all other commands default to a “Bypass” command. The OnCE Status register (OSR) is not visible when OnCE operation is disabled. In addition, OnCE Control register (OCR) functions are disabled, as is the operation of the *jd_de_b* input. Secure systems may choose to leave this signal negated until a security check has been performed. Other systems should tie this signal asserted to enable full OnCE operation. The *j_en_once_regssel* and *j_key_in* signals are provided to assist external logic performing security checks. Refer to [Section 6.3.15.15, “Enable Once Register Select \(*j_en_once_regssel*\),”](#) for a description of the *j_en_once_regssel* output signal, and to [Section 6.3.15.20, “Key Data In \(*j_key_in*\),”](#) for a description of the *j_key_in* input signal.

The *jd_en_once* input must only change state during the Test-Logic-Reset, Run-Test/Idle, or Update_DR TAP states. A new value takes effect after one additional *j_tclk* cycle of synchronization.

6.3.13.2 Debug Session (*jd_debug_b*)

The *jd_debug_b* active-low output signal is asserted when the processor first enters into debug mode. It remains asserted for the duration of a “debug session”.

NOTE

A debug session includes single-step operations (Go+NoExit OnCE commands). That is, *jd_debug_b* remains asserted during OnCE single-step executions.

This signal is provided to allow system resources to be aware that access is occurring for debug purposes, thus allowing certain resource side effects to be frozen or otherwise controlled. Examples might include FIFO state change control, control of side-effects of register or memory accesses, etc. Refer to [Section 8.4.4.3, “e200 OnCE Debug Output \(*jd_debug_b*\)](#),” for additional information on this signal.

6.3.13.3 Debug Request (*jd_de_b*)

This signal is the debug mode request input. This signal is not internally synchronized by the e200 core, thus it must meet setup and hold time constraints relative to *j_tclk*. To be recognized, it must be held asserted for a minimum of two *j_tclk* periods, and the *jd_en_once* input must be in the asserted state. *jd_de_b* is synchronized to *m_clk* in the debug module before being sent to the processor (two clocks).

This signal is normally the input from the top-level *DE_b* open-drain bidirectional I/O cell. Refer to [Section 8.4.4.2, “OnCE Debug Request/Event \(*jd_de_b*, *jd_de_en*\)](#),” for additional information on this signal.

6.3.13.4 DE_b Active High Output Enable (*jd_de_en*)

This output signal is an active-high enable for the top-level *DE_b* open-drain bidirectional I/O cell. This signal is asserted for three *j_tclk* periods upon processor entry into debug mode. Refer to [Section 8.4.4.2, “OnCE Debug Request/Event \(*jd_de_b*, *jd_de_en*\)](#),” for additional information on this signal.

6.3.13.5 Processor Clock On (*jd_mclk_on*)

This active-high input signal is driven by system level clock control logic to indicate that the processor’s *m_clk* input is active. This signal is synchronized to *j_tclk* and provided as a status bit in the OnCE Status register.

6.3.13.6 Watchpoint Events (*jd_watchpoint*[0:5])

The *jd_watchpoint*[0:5] active-high output signals are used to indicate that a watchpoint has occurred. Each debug address compare function (IAC1-4, DAC1-2) is capable of triggering a watchpoint output. Refer to [Section 8.5, “Watchpoint Support](#),” for the signal assignments of each watchpoint source.

6.3.14 Development Support (Nexus2+) Signals

The following interface signals are provided to assist in implementing a real-time development tool capability with a controller external to the e200 core. These signals are optional and are described in an external Nexus2/3 specification.

Table 6-16. e200 Development Support (Nexus2+) Signals

Signal	Type	Description
<i>nex_mcko</i>	O	Nexus Clock Output
<i>nex_rdy_b</i>	O	Nexus Ready Output
<i>nex_evto_b</i>	O	Nexus Event-Out Output
<i>nex_evti_b</i>	I	Nexus Event-In Input
<i>nex_mdo[n:0]</i>	O	Nexus Message Data Output
<i>nex_mseo_b[1:0]</i>	O	Nexus Message Start/End Output

6.3.15 JTAG Support Signals

Table 6-17 details the primary JTAG interface signals. These signals are usually connected directly to device pins (except for *j_tdo*, which needs tri-state and edge support logic). However, this may not be the case when JTAG TAP controllers are concatenated together.

Table 6-17. JTAG Primary Interface Signals

Signal Name	Type	Description
<i>j_trst_b</i>	I	JTAG test reset
<i>j_tclk</i>	I	JTAG test clock
<i>j_tms</i>	I	JTAG test mode select
<i>j_tdi</i>	I	JTAG test data input
<i>j_tdo</i>	O	Test data out to master controller or pad
<i>j_tdo_en</i> ¹	O	Enables TDO output buffer

¹ *j_tdo_en* is asserted when the TAP controller is in the *shift_dr* or *shift_ir* state.

6.3.15.1 JTAG/OnCE Serial Input (*j_tdi*)

Data and commands are provided to the OnCE controller through the *j_tdi* pin. Data is latched on the rising edge of the *j_tclk* serial clock. Data is shifted into the OnCE serial port least significant bit (LSB) first.

6.3.15.2 JTAG/OnCE Serial Clock (*j_tclk*)

The *j_tclk* pin supplies the serial clock to the OnCE control block. The serial clock provides pulses required to shift data and commands into and out of the OnCE serial port. (Data is clocked into the OnCE

on the rising edge and is clocked out of the OnCE serial port on the rising edge.) The debug serial clock frequency must be no greater than 50% of the processor clock frequency.

6.3.15.3 JTAG/OnCE Serial Output (*j_tdo*)

Serial data is read from the OnCE block through the *j_tdo* pin. Data is always shifted out the OnCE serial port least significant bit (LSB) first. When data is clocked out of the OnCE serial port, *j_tdo* changes on the rising edge of *j_tclk*. The *j_tdo* output signal is always driven.

An external system-level TDO pin may be tri-stateable and should be actively driven in the shift-IR and shift-DR controller states. The *j_tdo_en* signal is supplied to indicate when an external TDO pin should be enabled and is asserted during the shift-IR and shift-DR controller states. In addition, for IEEE Std 1149™ compatibility, the system level pin should change state on the falling edge of TCLK.

6.3.15.4 JTAG/OnCE Test Mode Select (*j_tms*)

The *j_tms* input is used to cycle through states in the OnCE Debug Controller. Toggling the *j_tms* pin while clocking with *j_tclk* controls transitions through the TAP state controller.

6.3.15.5 JTAG/OnCE Test Reset (*j_trst_b*)

The *j_trst_b* input is used to externally reset the OnCE controller by placing it in the Test-Logic-Reset state.

[Table 6-18](#) details additional signals which may be used to support external JTAG data registers using the e200 TAP controller.

Table 6-18. JTAG Signals Used to Support External Registers

Signal Name	Type	Description
<i>j_tst_log_rst</i>	O	Indicates the TAP controller is in the Test-Logic-Reset state
<i>j_rti</i>	O	JTAG controller run-test/idle state
<i>j_capture_ir</i>	O	Indicates the TAP controller is in the capture IR state
<i>j_shift_ir</i>	O	Indicates the TAP controller is in shift IR state
<i>j_update_ir</i>	O	Indicates the TAP controller is in update IR state
<i>j_capture_dr</i>	O	Indicates the TAP controller is in the capture DR state
<i>j_shift_dr</i>	O	Indicates the TAP controller is in shift DR state
<i>j_update_gp_reg</i>	O	Updates JTAG controller general-purpose data register
<i>j_gp_regsel[0:11]</i>	O	General-purpose external JTAG register select
<i>j_en_once_regsel</i>	O	External Enable OnCE register select
<i>j_key_in</i>	I	Serial data from external key logic
<i>j_lsrl_regsel</i>	O	External LSRL register select
<i>j_serial_data</i>	I	Serial data from external JTAG register(s)

6.3.15.6 Test-Logic-Reset (*j_tst_log_rst*)

This signal indicates the TAP controller is in the Test-Logic-Reset state.

6.3.15.7 Run-Test/Idle (*j_rti*)

This signal indicates the TAP controller is in the Run-Test/Idle state.

6.3.15.8 Capture IR (*j_capture_ir*)

This signal indicates the TAP controller is in the Capture_IR state.

6.3.15.9 Shift IR (*j_shift_ir*)

This signal indicates the TAP controller is in the Shift_IR state.

6.3.15.10 Update IR (*j_update_ir*)

This signal indicates the TAP controller is in the Update_IR state.

6.3.15.11 Capture DR (*j_capture_dr*)

This signal indicates the TAP controller is in the Capture_DR state.

6.3.15.12 Shift DR (*j_shift_dr*)

This signal indicates the TAP controller is in the Shift_DR state.

6.3.15.13 Update DR (*j_update_gp_reg*)

This signal indicates the TAP controller is in the Update_DR state and that the R/W bit in the OnCE Command register is low (write command). The *j_gp_regsel[0:11]* signals should be monitored to see which register, if any, needs to be updated.

6.3.15.14 Register Select (*j_gp_regsel*)

The outputs shown in [Table 6-19](#) are a decode of the REGSEL[0:6] field in the OnCE Command Register (OCMD). They are used to specify which external general purpose JTAG register to access via the e200 TAP controller.

Table 6-19. JTAG General Purpose Register Select Decoding

Signal Name	Type	Description
<i>j_gp_regsel[0]</i>	O	REGSEL[0:6]=7'h70
<i>j_gp_regsel[1]</i>	O	REGSEL[0:6]=7'h71
<i>j_gp_regsel[2]</i>	O	REGSEL[0:6]=7'h72
<i>j_gp_regsel[3]</i>	O	REGSEL[0:6]=7'h73

Table 6-19. JTAG General Purpose Register Select Decoding (continued)

Signal Name	Type	Description
<i>j_gp_regsel</i> [4]	O	REGSEL[0:6]=7'h74
<i>j_gp_regsel</i> [5]	O	REGSEL[0:6]=7'h75
<i>j_gp_regsel</i> [6]	O	REGSEL[0:6]=7'h76
<i>j_gp_regsel</i> [7]	O	REGSEL[0:6]=7'h77
<i>j_gp_regsel</i> [8]	O	REGSEL[0:6]=7'h78
<i>j_gp_regsel</i> [9]	O	REGSEL[0:6]=7'h79
<i>j_gp_regsel</i> [10]	O	REGSEL[0:6]=7'h7A
<i>j_gp_regsel</i> [11]	O	REGSEL[0:6]=7'h7B

6.3.15.15 Enable Once Register Select (*j_en_once_regsel*)

The *j_en_once_regsel* output is asserted when a decode of the REGSEL[0:6] field in the OnCE Command Register (OCMD) indicates an external Enable_OnCE register is selected (0b111110 encoding) for access via the e200 TAP controller. This control signal may be used by external security logic to assist in controlling the *jd_enable_once* input signal. The external Enable_OnCE register should be muxed onto the *j_serial_data* input (Refer to [Section 6.3.15.19, “Serial Data \(*j_serial_data*\)”](#)). During the Shift_DR state, *j_serial_data* is supplied to the *j_tdo* output.

6.3.15.16 External Nexus Register Select (*j_nexus_regsel*)

The *j_nexus_regsel* output is asserted when a decode of the REGSEL[0:6] field in the OnCE Command Register (OCMD) indicates an external Nexus register is selected (0b1111100 encoding) for access via the e200 TAP controller.

6.3.15.17 External Shared Nexus Control Register Select (*j_sncr_regsel*)

The *j_sncr_regsel* output is asserted when a decode of the REGSEL[0:6] field in the OnCE Command Register (OCMD) indicates an external Shared Nexus Control register is selected (0b1101111 encoding) for access via the e200 TAP controller.


6.3.15.18 External LSRL Register Select (*j_lsrl_regsel*)

The *j_lsrl_regsel* output is asserted when a decode of the REGSEL[0:6] field in the OnCE Command Register (OCMD) indicates an external LSRL register is selected (0b1111101 encoding) for access via the e200 TAP controller.

6.3.15.19 Serial Data (*j_serial_data*)

This input signal receives serial data from external JTAG registers. All external registers share this one serial output back to the core, therefore it must be muxed using the *j_gp_regsel*[0:11], *j_lsrl_regsel*, and *j_en_once_regsel* signals. The data is internally routed to *j_tdo*.

Figure 6-3 shows one example of how an external JTAG register set (2) could be designed using the inputs and outputs provided and by the JTAG primary inputs themselves. The main components are a clock generation unit, a JTAG shifter (load, shift, hold, clr), the registers (load, hold, clr), and an input mux to the shifter for the serial output back to the e200 core. The shifter and the registers may be as wide as the application warrants [0:x]. The length determines the number of states the TAP controller is held in Shift_DR (x+1).



NOTES:

1. $\text{clk_shfter} = \text{j_tclk} \text{ and } (\text{j_shift_dr} \mid \text{j_capture_dr})$
2. $\text{clk_reg0} = \text{j_tclk} \text{ and } \text{j_update_gp_reg} \text{ and } \text{j_gp_regssel}[0]$
3. $\text{clk_reg1} = \text{j_tclk} \text{ and } \text{j_update_gp_reg} \text{ and } \text{j_gp_regssel}[1]$

Figure 6-3. Example External JTAG Register Design

6.3.15.20 Key Data In (*j_key_in*)

This input signal receives serial data from logic to indicate a key or other value to be scanned out in the Shift_IR state when the current value in the IR is the Enable_OnCE instruction. This input is provided to assist in implementing security logic outside of the processor which conditionally asserts *jd_en_once*. During the Shift_IR state, when *jd_en_once* is negated, this input is sampled on the rising edge of *j_tclk*, and after a two clock delay the data is internally routed to *j_tdo*. This allows provision of a key value via the *j_tdo* output following a transition from Capture_IR to Shift_IR. The key value is provided via the *j_key_in* input.

6.3.16 JTAG ID Signals

Table 6-20 shows the JTAG ID register unique to Freescale as specified by the *IEEE 1149.1 JTAG Specification*. Note that bit 31 is the MSB of this register.

Table 6-20. JTAG Register ID Fields

Bit Field	Type	Description	Value
[31:28]	Variable	Version Number	Variable
[27:22]	Fixed	Design Center Number	6'b011111
[21:12]	Variable	Sequence Number	Variable
[11:1]	Fixed	Freescale Manufacturer ID	11'b00000001110
0	Fixed	JTAG ID Register Identification Bit	1'b1

The e200 core shifts out a “1” as the first bit on *j_tdo* if the Shift_DR state is entered directly from the test-logic-reset state. This is per the JTAG specification and informs any JTAG controller that an ID register exists on the part. The e200 JTAG ID register is accessed by writing the OCMR (OnCE Command Register) with the value 7'h02 in the REGSEL[0:6] field.

The JTAG ID bit, manufacturer ID field, and design center number are fixed by the JTAG Consortium and/or Freescale. The version numbers and the two most significant bits (MSBs) of the sequence number are variable and brought out to external ports. The lower eight bits of the sequence number are variable and strapped internally to track variations in processor deliverables.

Table 6-21 shows the inputs to the JTAG ID register that are input ports on the e200 core. These bits are provided for a customer to track revisions of a device using the e200 core.

Table 6-21. JTAG ID Register Inputs

Signal Name	Type	Description
<i>j_id_sequence</i> [0:1]	I	JTAG ID register (2 MSBs of sequence field)
<i>j_id_version</i> [0:3]	I	JTAG ID register version field

6.3.16.1 JTAG ID Sequence (*j_id_sequence*[0:1])

The *j_id_sequence*[0:1] inputs correspond to the two MSBs of the 10-bit sequence number in the JTAG ID register. These inputs are normally static. They are provided for the customer for further component variation identification.

6.3.16.2 JTAG ID Sequence (*j_id_sequence*[2:9])

The *j_id_sequence*[2:9] field is internally strapped to track variations in processor and module deliverables. Each e200 deliverable has a unique sequence number. Additionally, each revision of these modules can be identified by unique sequence numbers.

6.3.16.3 JTAG ID Version (*j_id_version*[0:3])

The *j_id_version*[0:3] inputs correspond to the 4-bit version number in the JTAG ID register. These inputs are normally static. They are provided to the customer for strapping in order to facilitate easy identification of component variants.

6.4 Timing Diagrams

6.4.1 Processor Instruction/Data Transfers

Transfer of data between the core and peripherals involves the address bus, data busses, and control and attribute signals. The address and data buses are parallel, non-multiplexed buses, supporting byte, halfword, three byte, and word transfers. All bus input and output signals are sampled and driven with respect to the rising edge of the *m_clk* signal. The core moves data on the bus by issuing control signals and using a handshake protocol to ensure correct data movement.

The memory interface operates in a pipelined fashion to allow additional access time for memory and peripherals. AHB transfers consist of an address phase which lasts only a single cycle, followed by the data phase which may last for one or more cycles depending on the state of the *p_hready* signal.

Read transfers consist of a request cycle, where address and attributes are driven along with a transfer request, and one or more memory access cycles to perform accesses and return data to the CPU for alignment, sign or zero extension, and forwarding.

Write transfers consist of a request cycle, where address and attributes are driven along with a transfer request, and one or more data drive cycles where write data is driven and external devices accept write data for the access.

Access requests are generated in an overlapped fashion in order to support sustained single cycle transfers. Up to two access requests may be in progress at any one cycle, one access outstanding and a second in the pending request phase.

Access requests are assumed to be accepted as long as there are no accesses in progress, or if an access in progress is terminated during the same cycle a new request is active (*p_hready* asserted). Once an access has been accepted, the BIU is free to change the current request at any time, even if part of a burst transfer.

The local memory control logic is responsible for proper pipelining and latching of all interface signals to initiate memory accesses.

The system hardware can use the *p_hresp*[2:0] signals to signal that the current bus cycle has an error when a fault is detected, using the ERROR response encoding. ERROR assertion requires a two cycle response. In the first cycle of the response, the *p_hresp*[2:0] signals are driven to indicate ERROR and *p_hready* must be negated. During the following cycle, the ERROR response must continue to be driven, and *p_hready* must be asserted. When the core recognizes a bus error condition for an access at the end of the first cycle of the two cycle error response, a subsequent pending access request may be removed by the BIU driving the *p_htrans*[2:0] signals to the IDLE state in the second cycle of the two cycle error response. Not all pending requests are removed, however.

When a bus cycle is terminated with a bus error, the core can enter storage error exception processing immediately following the bus cycle, or it can defer processing the exception.

The instruction prefetch mechanism requests instruction words from the instruction memory unit before it is ready to execute them. If a bus error occurs on an instruction fetch, the core does not take the exception until it attempts to use the instruction. Should an intervening instruction cause a branch, or should a task switch occur, the storage error exception for the unused access does not occur. A bus error termination for any write access or read access that reference data specifically requested by the execution unit causes the core to begin exception processing.

6.4.1.1 Basic Read Transfer Cycles

During a read transfer, the core receives data from a memory or peripheral device. Figure 6-4 illustrates functional timing for basic read transfers. Clock-by-clock descriptions of activity in Figure 6-4 follow.




Figure 6-4. Basic Read Transfers

Clock 1 (C1):

The first read transfer starts in clock cycle 1. During C1, the core places valid values on the address bus and transfer attributes. The burst type (*p_hburst*[2:0]), protection control (*p_hprot*[5:0]), and transfer type (*p_htrans*[1:0]) attributes identify the specific access type. The transfer size attributes (*p_hsize*[1:0])

indicates the size of the transfer. The byte strobes ($p_hbstrb[3:0]$) are driven to indicate active byte lanes. The write (p_hwrite) signal is driven low for a read cycle.

The core asserts transfer request ($p_htrans = \text{NONSEQ}$) during C1 to indicate that a transfer is being requested. Because the bus is currently idle, (0 transfers outstanding), the first read request to addr_x is considered *taken* at the end of C1. The default slave drives an ready/OKAY response for the current idle cycle.

Clock 2 (C2):

During C2, the addr_x memory access takes place using the address and attribute values which were driven during C1 to enable reading of one or more bytes of memory. Read data from the slave device is provided on the p_hrdata inputs. The slave device responds by asserting p_hready to indicate the cycle is completing and drives an OKAY response.

Another read transfer request is made during C2 to addr_y ($p_htrans = \text{NONSEQ}$), and because the access to addr_x is completing, it is considered *taken* at the end of C2.

Clock 3 (C3):

During C3, the addr_y memory access takes place using the address and attribute values which were driven during C2 to enable reading of one or more bytes of memory. Read data from the slave device for addr_y is provided on the p_hrdata inputs. The slave device responds by asserting p_hready to indicate the cycle is completing and drives an OKAY response.

Another read transfer request is made during C3 to addr_z ($p_htrans = \text{NONSEQ}$), and because the access to addr_y is completing, it is considered *taken* at the end of C3.

Clock 4 (C4):

During C4, the addr_z memory access takes place using the address and attribute values which were driven during C3 to enable reading of one or more bytes of memory. Read data from the slave device for addr_z is provided on the p_hrdata inputs. The slave device responds by asserting p_hready to indicate the cycle is completing and drives an OKAY response.

The CPU has no more outstanding requests, so p_htrans indicates IDLE. The address and attribute signals are thus undefined.

6.4.1.2 Read Transfer with Wait State

Figure 6-5 shows an example of wait state operation. Signal p_hready for the first request (addr_x) is not asserted during C2, so a wait state is inserted until p_hready is recognized (during C3).

Meanwhile, a subsequent request has been generated by the CPU for addr_y which is not *taken* in C2, because the previous transaction is still outstanding. The address and transfer attributes remain driven in cycle C3 and are taken at the end of C3 because the previous access is completing. Data for addr_x and a ready/OKAY response is driven back by the slave device. In cycle C4, a request for addr_z is made. The request for access to addr_z is taken at the end of C4, and during C5, the data and a ready/OKAY response is provided by the slave device. In cycle C5, no further accesses are requested.




Figure 6-5. Read Transfer with Wait-state

6.4.1.3 Basic Write Transfer Cycles

During a write transfer, the core provides write data to a memory or peripheral device. Figure 6-6 illustrates functional timing for basic write transfers. Clock-by-clock descriptions of activity in Figure 6-6 follow.




Figure 6-6. Basic Write Transfers

Clock 1 (C1):

The first write transfer starts in clock cycle 1. During C1, the core places valid values on the address bus and transfer attributes. The burst type (*p_hburst*[2:0]), protection control (*p_hprot*[5:0]), and transfer type (*p_htrans*[1:0]) attributes identify the specific access type. The transfer size attributes (*p_hsize*[1:0]) indicates the size of the transfer. The byte strobes (*p_hbstrb*[3:0]) are driven to indicate active byte lanes. The write (*p_hwrite*) signal is driven high for a write cycle.

The core asserts transfer request (*p_htrans*= NONSEQ) during C1 to indicate that a transfer is being requested. Because the bus is currently idle, (0 transfers outstanding), the first read request to *addr_x* is considered *taken* at the end of C1. The default slave drives an ready/OKAY response for the current idle cycle.

Clock 2 (C2):

During C2, the write data for the access is driven, and the *addr_x* memory access takes place using the address and attribute values which were driven during C1 to enable writing of one or more bytes of

memory. The slave device responds by asserting p_hready to indicate the cycle is completing and drives an OKAY response.

Another write transfer request is made during C2 to $addr_y$ ($p_htrans = NONSEQ$), and because the access to $addr_x$ is completing, it is considered *taken* at the end of C2.

Clock 3 (C3):

During C3, write data for $addr_y$ is driven, and the $addr_y$ memory access takes place using the address and attribute values which were driven during C2 to enable writing of one or more bytes of memory. The slave device responds by asserting p_hready to indicate the cycle is completing and drives an OKAY response.

Another write transfer request is made during C3 to $addr_z$ ($p_htrans = NONSEQ$), and because the access to $addr_y$ is completing, it is considered *taken* at the end of C3.

Clock 4 (C4):

During C4, write data for $addr_z$ is driven, and the $addr_z$ memory access takes place using the address and attribute values which were driven during C3 to enable reading of one or more bytes of memory. The slave device responds by asserting p_hready to indicate the cycle is completing and drives an OKAY response.

The CPU has no more outstanding requests, so p_htrans indicates IDLE. The address and attribute signals are thus undefined.

6.4.1.4 Write Transfer with Wait States

Figure 6-7 shows an example of write wait state operation. Signal p_hready for the first request ($addr_x$) is not asserted during C2, so a wait state is inserted until p_hready is recognized (during C3).

Meanwhile, a subsequent request has been generated by the CPU for $addr_y$ which is not *taken* in C2, because the previous transaction is still outstanding. The address, transfer attributes, and write data remain driven in cycle C3 and are taken at the end of C3 because a ready/OKAY response is driven back by the slave device for the previous access. In cycle C4, a request for $addr_z$ is made. The request for access to $addr_z$ is taken at the end of C4, and during C5, a ready/OKAY response is provided by the slave device. In cycle C5, no further accesses are requested.




Figure 6-7. Write Transfer with Wait-State

6.4.1.5 Read and Write Transfers

Figure 6-8 shows a sequence of read and write cycles.




Figure 6-8. Single Cycle Read and Write Transfers

The first read request (addr_x) is *taken* at the end of cycle C1 because the bus is idle.

The second read request (addr_y) is *taken* at the end of C2 because a ready/OKAY response is asserted during C2 for the first read access (addr_x). During C3, a request is generated for a write to addr_y , which is taken at the end of C3 because the second access is terminating.

Data for the addr_z write cycle is driven in C4, the cycle after the access is *taken*, and a ready/OKAY response is signaled to complete the write cycle to addr_z .

Figure 6-9 shows another sequence of read and write cycles. This example shows an interleaved write access between two reads.




Figure 6-9. Single Cycle Read and Write Transfers—2

The first read request (addr_x) is *taken* at the end of cycle C1 because the bus is idle.

The first write request (addr_y) is *taken* at the end of C2 because the first access is terminating (addr_x).

Data for the addr_y write cycle is driven in C3, the cycle after the access is *taken*. Also during C3, a request is generated for a read to addr_z, which is taken at the end of C3 because the write access is terminating.

During C4, the addr_y write access is terminated, and no further access is requested

Figure 6-10 shows another sequence of read and write cycles. In this example, reads incur a single wait state.




Figure 6-10. Multi-Cycle Read and Write Transfers

The first read request (addr_x) is *taken* at the end of cycle C1 because the bus is idle.

The second read request (addr_y) is not *taken* at the end of cycle C2 because no ready response is signaled and only one access can be outstanding (addr_x). It is taken at the end of C3 once the first read request has signaled a ready/OKAY response.

The first write request (addr_z) is not taken during C4 because a ready response is not asserted during C4 for the second read access (addr_y). During C5, the request for a write to addr_z is taken because the second access is terminating.

Data for the addr_z write cycle is driven in C6, the cycle after the access is *taken*.

During C6, the addr_z write access is terminated and the addr_w write request is *taken*.

During C7, data for the addr_w write access is driven, and a ready/OKAY response is asserted to complete the write cycle to addr_w.

Figure 6-11 shows another sequence of read and write cycles. In this example, reads incur a single wait state.




Figure 6-11. Multi-Cycle Read and Write Transfers—2

The first read request (addr_x) is *taken* at the end of cycle C1 because the bus is idle.

The first write request (addr_y) is not *taken* at the end of cycle C2 because no ready response is signaled and only one access can be outstanding (addr_x). It is taken at the end of C3 once the first read request has signaled a ready/OKAY response.

Data for the addr_y write cycle is driven in C4, the cycle after the access is *taken*.

The second read request (addr_z) is taken during C4 because the addr_y write is terminating.

A second write request (addr_w) is not taken at the end of C5 because the second read access is not terminating, thus it continues to drive the address and attributes into cycle C6.

During C6, the addr_z read access is terminated and the addr_w write access is taken.

In cycle C7, data for the addr_w write access is driven. During C7, a ready/OKAY response is asserted to complete the write cycle to addr_w . No further accesses are requested, so p_htrans signals IDLE.

6.4.1.6 Misaligned Accesses

Figure 6-12 illustrates functional timing for a misaligned read transfer. The read to addr_x is misaligned across a 32-bit boundary.




Figure 6-12. Misaligned Read Transfer

The first portion of the misaligned read transfer starts in C1. During C1, the core places valid values on the address bus and transfer attributes. The p_hwrite signal is driven low for a read cycle. The transfer size attributes (p_hsize) indicate the size of the transfer. Even though the transfer is misaligned, the size value driven corresponds to the size of the entire misaligned data item. $p_hunalign$ is driven high to indicate that the access is misaligned. The p_hbstrb outputs are asserted to indicate the active byte lanes for the read, which may not correspond to size and low-order address outputs. p_htrans is driven to NONSEQ.

During C2, the addr_x memory access takes place using the address and attribute values which were driven during C1 to enable reading of one or more bytes of memory.

The second portion of the misaligned read transfer request is made during C2 to addr_{x+} (which is aligned to the next higher 32-bit boundary), and because the first portion of the misaligned access is completing, it is *taken* at the end of C2. The p_htrans signals indicate NONSEQ. The size value driven is the size of the remaining bytes of data in the misaligned read, rounded up (for the 3-byte case) to the next higher power-of-2. The p_hbstrb signals indicate the active byte lanes. For the second portion of a misaligned

transfer, the *p_hunalign* signal is driven high for the 3-byte case (low for all others). The next read access is requested in C3 and *p_htrans* indicates NONSEQ. *p_hunalign* is negated, because this access is aligned.

Figure 6-13 illustrates functional timing for a misaligned write transfer. The write to addr_x is misaligned across a 32-bit boundary.



Figure 6-13. Misaligned Write Transfer

The first portion of the misaligned write transfer starts in C1. During C1, the core places valid values on the address bus and transfer attributes. The *p_hwrite* signal is driven high for a write cycle. The transfer size attribute (*p_hsize*) indicate the size of the transfer. Even though the transfer is misaligned, the size value driven corresponds to the size of the entire misaligned data item. *p_hunalign* is driven high to indicate that the access is misaligned. The *p_hbstrb* outputs are asserted to indicate the active byte lanes for the write, which may not correspond to size and low-order address outputs. *p_htrans* is driven to NONSEQ.

During C2, data for addr_x is driven, and the addr_x memory access takes place using the address and attribute values which were driven during C1 to enable writing of one or more bytes of memory.

The second portion of the misaligned write transfer request is made during C2 to addr_{x+} (which is aligned to the next higher 32-bit boundary), and because the first portion of the misaligned access is completing, it is *taken* at the end of C2. The *p_htrans* signals indicate NONSEQ. The size value driven is the size of the remaining bytes of data in the misaligned write, rounded up (for the 3-byte case) to the next higher

power-of-2. The *p_hbstrb* signals indicate the active byte lanes. For the second portion of a misaligned transfer, the *p_hunalign* signal is driven high for the 3-byte case (low for all others).

The next write access is requested in C3 and *p_htrans* indicates NONSEQ. *p_hunalign* is negated, because this access is aligned.

An example of a misaligned write cycle followed by an aligned read cycle is shown in [Figure 6-14](#). It is similar to the previous example in [Figure 6-13](#).




Figure 6-14. Misaligned Write, Single Cycle Read Transfer

6.4.1.7 Burst Accesses

Figure 6-15 illustrates functional timing for a burst read transfer.




Figure 6-15. Burst Read Transfer

The *p_hburst* signals indicate INCR for all burst transfers. The *p_hunalign* signal is negated. *p_hsize* indicates 32-bits, and all four *p_hbstrb* signals are asserted. The burst address is aligned to a 32-bit boundary and increments by words. Note that in this example four beats are shown, but in operation the burst may be of any length including only a single beat.

NOTE

Bursts may be interrupted immediately at any time, and be followed by any type of cycle. No idle cycle is required.

Figure 6-16 illustrates functional timing for a burst read with wait-state transfer.




Figure 6-16. Burst Read with Wait-State Transfer

The first cycle of the burst incurs a single wait-state.

Figure 6-17 illustrates functional timing for a burst write transfer.




Figure 6-17. Burst Write Transfer

Figure 6-16 illustrates functional timing for a burst write with wait-state transfer.




Figure 6-18. Burst Write with Wait-State Transfer

The first cycle of the burst incurs a single wait-state. Data for the second beat of the burst is valid the cycle after the second beat is *taken*.

Figure 6-19 illustrates functional timing for a pair of burst read transfers.




Figure 6-19. Burst Read Transfers

Note that in this example the first burst is two beats long and is followed immediately by a second burst, which is unrelated to the first.

NOTE

Bursts may be of any length (including a single beat) and may be followed immediately by any type of transfer. No idle cycles are required.

Figure 6-20 illustrates functional timing for a burst read with wait-state transfer where the second beat to addr x+ is retracted and replaced with a new burst transfer.




Figure 6-20. Burst Read with Wait-State Transfer, Retraction

The first cycle of the burst incurs a single wait-state, and the burst is replaced by another burst.

Figure 6-21 illustrates functional timing for a burst write transfer. The second burst is only one beat long.




Figure 6-21. Burst Write Transfers, Single Beat Burst

This same scenario can occur for read bursts as well.

6.4.1.8 Address Retraction

Address retraction is the process of replacing an existing request with a new request unrelated to the first request. Although the AMBA AHB protocol requires an access request to remain driven unchanged once presented on the bus, higher system performance may be obtained if this aspect of the protocol is modified to allow an access request to be changed prior to being taken. e200z1Zen Z0n2p and Zen Z0Hn2p always performs address retraction under conditions in which performance may be optimized. Figure 6-22 shows an example of address retraction during wait state operation. Signal *p_hready* for the first request (*addr_x*) is not asserted during C2, so a wait state is inserted until *p_hready* is recognized (during C3).

Meanwhile, a subsequent request has been generated by the CPU for *addr_y* which is not *taken* in C2, because the previous transaction is still outstanding. The address and transfer attributes are retracted in cycle C3, and a new access request to *addr_z* is requested and are taken at the end of C3 because the previous access is completing. Data for *addr_x* and a ready/OKAY response is driven back by the slave device. In cycle C4, a request for *addr_w* is made. The request for access to *addr_w* is taken at the end of C4, and during

C5, the data and a ready/OKAY response is provided by the slave device. In cycle C5, no further accesses are requested.




Figure 6-22. Read Transfer with Wait-State, Address Retraction

Figure 6-23 illustrates functional timing for a burst read with wait-state transfer where the second beat to addr x+ is retracted and replaced with a new burst transfer.




Figure 6-23. Burst Read with Wait-State Transfer, Retraction

The first cycle of the burst incurs a single wait-state, and the second beat of the burst driven in C2 burst is replaced by another burst in C3. Replacement by a single access is also possible.

Address retraction does not occur on a requested write cycle, only on read cycles, and may occur any time during a burst cycle as well.

6.4.1.9 Error Termination Operation

The *p_hresp*[2:0] inputs are used to signal an error termination for an access in progress. The ERROR encoding is used in conjunction with the assertion of *p_hready* to terminate a cycle with error. Error termination is a two-cycle termination; the first cycle consists of signaling the ERROR response on *p_hresp*[2:0] while holding *p_hready* negated, and during the second cycle, asserting *p_hready* while continuing to drive the ERROR response on *p_hresp*[2:0]. This two cycle termination allows the BIU to retract a pending access if it desires to do so. *p_htrans* may be driven to IDLE during the second cycle of the two-cycle error response, or may change to any other value, and a new access unrelated to the pending access may be requested. The cycle which may have been previously pending while waiting for a response

which terminates with error may be changed. It is not required to remain unchanged when an error response is received.

Figure 6-24 shows an example of error termination.




Figure 6-24. Read and Write Transfers, Instr. Read Error Termination

The first read request (addr_x) is *taken* at the end of cycle C1 because the bus is idle. It is an instruction prefetch.

The second read request (addr_y) is not *taken* at the end of C2 because the first access is still outstanding (no p_hready assertion). An error response is signaled by the addressed slave for addr_x by driving ERROR onto the $p_hresp[2:0]$ inputs. This is the first cycle of the two cycle error response protocol.

p_hready is asserted during C3 for the first read access (addr_x) while the ERROR encoding remains driven on $p_hresp[2:0]$, terminating the access. The read data bus is undefined.

In this example of error termination, the CPU continues to request an access to addr_y . It is taken at the end of C3. During C4, read data is supplied for the addr_y read, and the access is terminated normally during C4.

Also during C4, a request is generated for a read to addr_z , which is taken at the end of C4 because the second access is terminating.

Data for the addr_z read cycle is provided in C5, the cycle after the access is *taken*.

During C5, a ready/OKAY response is signaled to complete the read cycle to addr_z .

In this example of error termination, a subsequent access remained requested. This does not always occur when certain types of transfers are terminated with error. The following figures outline cases where an error termination for a given cycle causes a pending request to be aborted prior to initiation.

Figure 6-25 shows another example of error termination.

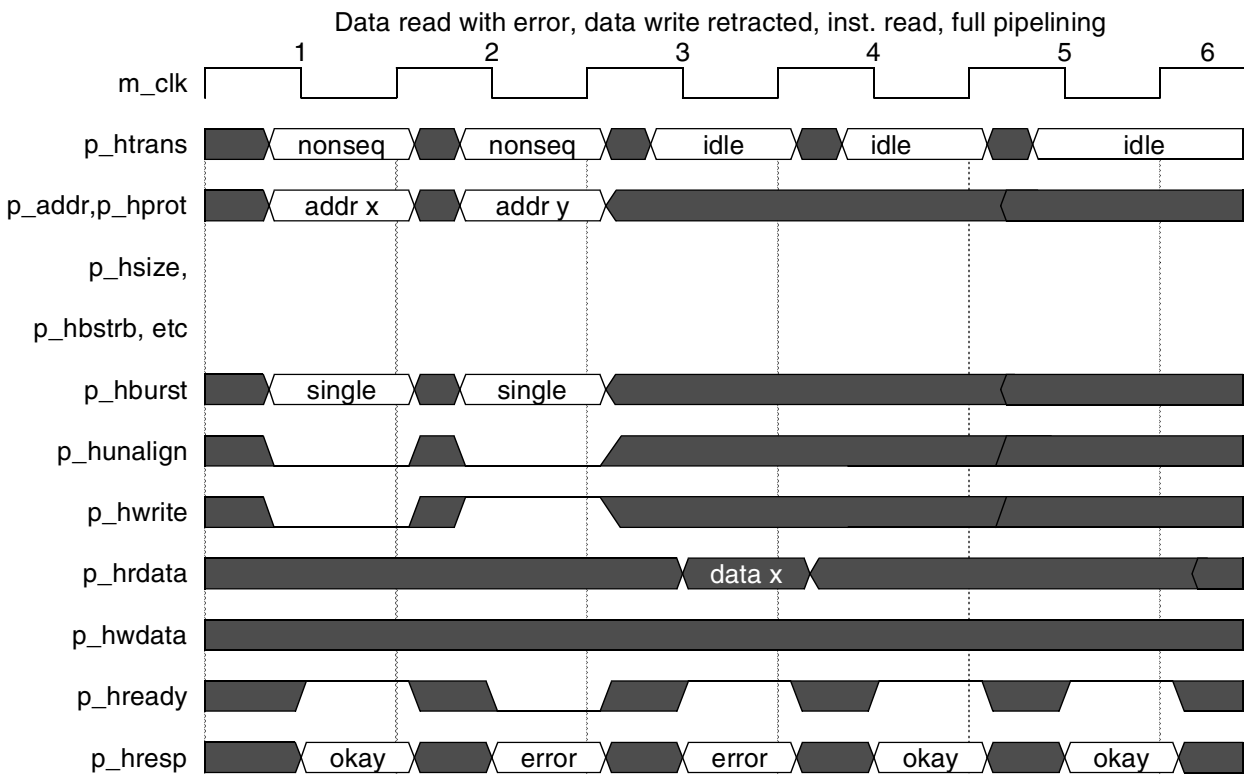


Figure 6-25. Data Read Error Termination

The first read request (addr_x) is *taken* at the end of cycle C1 because the bus is idle. It is a data read.

The second request (write to addr_y) is not *taken* at the end of C2 because the first access is still outstanding (no p_hready assertion). An error response is signaled by the addressed slave for addr_x by driving ERROR onto the $p_hresp[2:0]$ inputs. This is the first cycle of the two cycle error response protocol.

p_hready is asserted during C3 for the first read access (addr_x) while the ERROR encoding remains driven on $p_hresp[2:0]$, terminating the access. The read data bus is undefined.

In this example of error termination, the CPU retracts the requested access to addr_y by driving the p_htrans signals to the IDLE state during the second cycle of the two-cycle error response.

In this example of error termination, a subsequent access was aborted.

Figure 6-26 shows another example of error termination, this time on the initial portion of a misaligned write.




Figure 6-26. Misaligned Write Error Termination

The first portion of the misaligned write request is terminated with error. The second portion is aborted by the CPU during the second cycle of the two cycle error response.

6.4.2 Power Management

The following diagram shows the relationship of the wakeup control signal *p_wakeup* to the relevant input signals.




Figure 6-27. Wakeup Control Signal (*p_wakeup*)

6.4.3 Interrupt Interface

The following diagram shows the relationship of the interrupt input signals to the CPU clock. The *p_avec_b*, *p_extint_b*, *p_critint_b* and *p_voffset[0:9]* inputs as well as the *p_nmi_b* input must meet setup and hold timing relative to the rising edge of the *m_clk*. In addition, during each clock cycle in which either of the interrupt request inputs *p_extint_b* or *p_critint_b* are asserted, *p_avec_b* and *p_voffset[0:9]* are required to be in a valid state for the highest priority unmasked interrupt being requested.




Figure 6-28. Interrupt Interface Input Signals

Figure 6-29 and Figure 6-30 shows the relationship of the interrupt pending signal to the interrupt request inputs. Note that p_ipend is asserted combinationally from the p_extint_b , and $p_critint_b$, and p_nmi_b inputs.




Figure 6-29. e200 Interrupt Pending Operation




Figure 6-30. e200z0h Interrupt Pending Operation

Figure 6-31 shows the relationship of the interrupt acknowledge signal to the interrupt request inputs and exception vector fetching.




Figure 6-31. Interrupt Acknowledge Operation

In this example, an external input interrupt is requested in cycle 1. The *p_voffset*[0:9] inputs are driven with the vector offset for 'A', and *p_avec_b* is negated, indicating vectoring is desired. For this example, the bus is idle at the time of assertion. The CPU may sample a requested interrupt as early as the cycle it is initially requested, and does so in this example. The interrupt request and the vector offset and autovector input are sampled at the end of cycle 1. In cycle 3, the interrupt is acknowledged by the assertion of the *p_iack* output, indicating that the values present on interrupt inputs at the beginning of cycle 2 have been internally latched and committed to for servicing. Note that the interrupt vector lines have changed to a value of 'B' during cycle 2, and the *p_critint_b* input has been asserted by the interrupt controller. The vector number / autovector signals must be consistent with the higher priority critical input request, thus must change at the same time the state of the interrupt request inputs change. Because the *p_iack* output asserts in cycle 3, it is indicating that the values present at the rise of cycle 2 (vector 'A') have been committed to. During cycle 3, the CPU begins instruction fetching of the handler for vector 'A'.

The new request for a subsequent critical interrupt ‘B’ was not received in time to be acted upon first. It is acknowledged after the fetch for the external input interrupt handler has been completed and has entered decode.

Note that the time between assertion of an interrupt request input and the acknowledgment of an interrupt may be multiple cycles, and the interrupt inputs may change during that interval. The CPU asserts the *p_iack* output to indicate the cycle at which an interrupt is committed to. In the following example, because the CPU was unable to acknowledge the external input interrupt during cycle 2 due to internal or external execution conditions, the critical input request was sampled. This case is shown in [Figure 6-32](#).




Figure 6-32. Interrupt Acknowledge Operation—2

6.4.4 Time Base Interface

The following figure shows the required relationships of the Time Base inputs. The electrical values associated with these timings may be found in the *Zen Integration Guide*.




Figure 6-33. Time Base Input Timing

6.4.5 JTAG Test Interface

The following figures show the relationships of the various JTAG related signals to the *j_tclk* input. The electrical values associated with these timings may be found in the *Zen Integration Guide*.




Figure 6-34. Test Clock Input Timing





Figure 6-35. *j_trst_b* Timing



TEST_ACC_PRT_TIM_01

Figure 6-36. Test Access Port Timing



Chapter 7 Power Management

7.1 Power Management

Power management is supported by e200 cores to minimize overall system power consumption. The e200 core provides the ability to initiate power management from external sources as well as through software techniques. The power states on the e200 core are described below.

7.1.1 Active State

The Active state is the default state for the e200 core in which all of its internal units are operating at full processor clock speed. In this state, the e200 core still provides dynamic power management in which individual internal functional units may stop clocking automatically whenever they are idle.

7.1.2 Waiting State

The e200 core enters the Waiting state as a result of executing a **wait** instruction. Following entry into the waiting state, instruction execution and bus activity is suspended. Most internal clocks are gated off in this state. The e200 core asserts *p_waiting* to indicate it is in the waiting state. Prior to entering the waiting state, all outstanding instructions and bus transactions are completed. The *m_clk* input should remain running while in the waiting state to allow for interrupt sampling, and to allow further transitions into the Halted or Stopped state if requested.

In the waiting state, the core is waiting for a valid unmasked pending interrupt request. Once a pending interrupt request is received, the core exits the waiting state and begin interrupt processing. The return program counter value points to the next instruction after the **wait** instruction. The interrupt can be an external input interrupt, various critical interrupts, a debug interrupt (based on ICMP), a non-maskable interrupt (on e200z0h), or a machine check interrupt (*p_mcp_b* assertion, etc.). Once the interrupt processing begins, the core does not return to the waiting state until another **wait** instruction is executed.

The waiting state can be temporarily exited and returned to if a request is made to enter hardware debug mode (various mechanisms), the Halted state, or the Stopped state. After exiting one of these states, the processor returns to the waiting state. While temporarily exited, the *p_waiting* output negates, and is re-asserted once the CPU returns to the waiting state.

7.1.3 Halted State

Instruction execution and bus activity is suspended in the Halted state. However, none of the internal clocks are gated off in this state. The e200 core asserts *p_halted* to indicate it is in the halted state. Prior to entering the halted state, all outstanding bus transactions are completed. The *m_clk* input should remain running while in the Halted state to allow further transitions into the Stopped state if requested.

7.1.4 Stopped State

The Stopped state is characterized as having all internal functional units of the e200 core stopped except the clock control state machine logic. The internal *m_clk* may be kept running to allow quick recovery to the full on state. Clocks are not running to functional units in this state. The Stopped state is reached after transitioning through the Halted state with the *p_stop* input asserted. The *p_stopped* output signal is asserted once the Powerdown state is reached.

While in the Stopped state, further power savings may be achieved by stopping the *m_clk* input. This is done externally by the system after the e200 core is safely in the Stopped state and has asserted the *p_stopped* output signal. To exit from the Stopped state, the system must first restart the *m_clk* input.

7.1.5 Power Management Pins

p_waiting—output pin asserted when the e200 core is in the Waiting state.

p_halt—input pin is asserted by system logic to request the core to go into the Halted state. Negating this pin causes the e200 core to transition back into the Active or Waiting state if *p_stop* is also negated.

p_halted—output pin asserted when the e200 core is in the Halted state.

p_stop—input pin is asserted by system logic to request that the e200 core go into the Powerdown state. Negating this pin causes the e200 core to transition back into the Halted state from the Stopped state.

p_stopped—output pin asserted when the e200 core is in the Stopped state.

p_doze, *p_nap*, and *p_sleep* output pins that reflects the state of *HID0[DOZE]*, *HID0[NAP]*., and *HID0[SLEEP]* respectively. These pins are qualified with *MSR[WE] = 1*. Interpretation of these signals is done by the system logic.

p_wakeup—output pin asserted when an interrupt is pending or other condition which requires the clock to be running.

7.1.6 Power Management Control Bits

The following bits are used by software to generate a request to enter a power-saving state and to choose the state to be entered:

- *MSR[WE]*—The *WE* bit is used to qualify assertion of the ***p_doze***, ***p_nap***, and ***p_sleep*** output pins to the system logic. When *MSR[WE]* is negated, these pins are negated. When *MSR[WE]* is set, these pins reflect the state of their respective control bits in the *HID0* register.
- *HID0[DOZE]*—The interpretation of the doze mode bit is done by the external system logic. Doze mode on the e200 core is intended to be the halted state with the clocks running.
- *HID0[NAP]*—The interpretation of the nap mode bit is done by the external system logic. Nap mode on the e200 core may be used for a powerdown state with the Time Base enabled.
- *HID0[SLEEP]*—The interpretation of the sleep mode bit is done by the external system logic. Sleep mode on the e200 core may be used for a powerdown state with the Time Base disabled.

7.1.7 Software Considerations for Power Management Using Wait Instructions

Executing a **wait** instruction causes the e200 core to complete instruction fetch and execution activity and await an interrupt. The *p_waiting* output is asserted once the Waiting state is entered. External system hardware may interpret the state of this signal and activate the *p_halt* and/or *p_stop* inputs to cause the e200 core to enter a quiescent state in which clocks may be disabled for low power operation.

Alternatively, system hardware may utilize some other clock control mechanism while the processor is in the Waiting state, and *p_wakeup* remains negated.

7.1.8 Software Considerations for Power Management Using Doze, Nap or Sleep

Setting MSR[WE] generates a request to enter a power saving state. The power saving state (doze, nap, or sleep) must be previously determined by setting the appropriate HID0 bit. Setting MSR[WE] has no direct effect on instruction execution, but it is simply reflected on *p_doze*, *p_nap*, and *p_sleep* depending on the setting of HID0[DOZE], HID0[NAP], and HID0[SLEEP] respectively. Note that the e200 core is not affected by assertion of these pins directly. External system hardware may interpret the state of these signals and activate the *p_halt* and/or *p_stop* inputs to cause the e200 core to enter a quiescent state in which clocks may be disabled for low power operation.

To ensure a clean transition into and out of a power saving mode, the following program sequence is recommended:

```
sync
mtmsr (WE)
isync
loop:br loop (optionally use a wait instruction)
```

An interrupt is typically used to exit a power saving state. The *p_wakeup* output is used to indicate to the system logic that an interrupt (or a debug request) has become pending. System logic uses this output to re-enable the clocks and exit a low power state. The interrupt handler is responsible for determining how to exit the low power branching loop if one is used. Wait instructions are exited automatically. The vectored interrupt capability provided by the core may be useful in assisting the determination if an external hardware interrupt is used to perform the wake-up.

7.1.9 Debug Considerations for Power Management

When a debug request is presented to the e200 core while in either the Waiting, Halted or Stopped state, the *p_wakeup* signal is asserted, and when *m_clk* is provided to the CPU, it temporarily exits the Waiting, Halted or Stopped state and enters Debug mode regardless of the assertion of *p_halt* or *p_stop*. The *p_waiting*, *p_halted* or *p_stopped* outputs are negated for the duration of the time the CPU remains in a debug session (*jd_debug_b* asserted). When the debug session is exited, the CPU re-samples the *p_halt* and *p_stop* inputs and re-enters the Halted or Stopped state as appropriate. If the CPU was previously waiting, and no interrupt was received while in the debug session, it re-enters the Waiting state and re-asserts *p_waiting*.



Chapter 8

Debug Support

This chapter describes the debug features of the e200 core.

8.1 Overview

Internal debug support in the e200 core allows for software and hardware debug by providing debug functions, such as instruction and data breakpoints and program trace modes. For software based debugging, debug facilities consisting of a set of software accessible debug registers and interrupt mechanisms are provided. These facilities are also available to a hardware based debugger which communicates using a modified IEEE 1149.1 Test Access Port (TAP) controller and pin interface. When hardware debug is enabled, the debug facilities are protected from software modification.

Software debug facilities are defined as part of Power Architecture Book E. e200 supports a subset of these defined facilities. In addition to the facilities defined in Power Architecture Book E, e200 provides additional flexibility and functionality in the form of linked instruction and data breakpoints, and sequential debug event detection. These features are also available to a hardware-based debugger.

The e200 core provides support for an external Nexus real-time debug module. Real-time debugging in a e200-based system is supported by a Nexus class 1 module.

8.1.1 Software Debug Facilities

e200 provides debug facilities to enable hardware and software debug functions, such as instruction and data breakpoints and program single stepping. The debug facilities consist of a set of debug control registers (DBCR0-2), a set of address compare registers (IAC1, IAC2, IAC3, IAC4, DAC1, and DAC2), a Debug Status Register (DBSR) for enabling and recording various kinds of debug events, and a special Debug interrupt type built into the interrupt mechanism (see [Section 5.7.10, “Debug Interrupt \(IVOR15\)”](#)). The debug facilities also provide a mechanism for software-controlled processor reset in a debug environment.

Software debug facilities are enabled by setting the internal debug mode bit in Debug Control register 0 (DBCR0[IDM]). When internal debug mode is enabled, debug events can occur, and can be enabled to record exceptions in the Debug Status register (DBSR). If enabled by MSR[DE], these recorded exceptions cause Debug interrupts to occur. When DBCR0[IDM] is cleared, (and DBCR0[EDM] is cleared as well), no debug events occur, and no status flags are set in DBSR unless already set. In addition, when DBCR0[IDM] is cleared (or is overridden by DBCR0[EDM] being set) no Debug interrupts occur, regardless of the contents of DBSR. A software Debug interrupt handler may access all system resources and perform necessary functions appropriate for system debug.

8.1.1.1 Power Architecture Book E Compatibility

The e200 core implements a subset of the Power Architecture Book E internal debug features. The following restrictions on functionality are present:

- Instruction address compares do not support compare on physical (real) addresses.
- Data address compares do not support compare on physical (real) addresses.
- Data value compares are not supported.

8.1.2 Additional Debug Facilities

In addition to the debug functionality defined in Power Architecture Book E, e200 provides capability to link instruction and data breakpoints, and also provides a sequential breakpoint control mechanism.

e200 also defines two new debug events (CIRPT, CRET) for debugging around critical interrupts.

In addition, e200 implements the Debug APU, which when enabled allows Debug Interrupts to utilize a dedicated set of save/restore registers (DSRR0, DSRR1) for saving state information when a Debug Interrupt occurs, and for restoring this state information at the end of a debug interrupt handler by means of the `se_rfdi` instruction.

8.1.3 Hardware Debug Facilities

The e200 core contains facilities that allow for external test and debugging. A modified IEEE 1149.1 control interface is used to communicate with the core resources. This interface is implemented through a standard 1149.1 TAP (test access port) controller.

By using public instructions, the external debugger can freeze or halt the e200 core, read and write internal state and debug facilities, single-step instructions, and resume normal execution.

Hardware Debug is enabled by setting the External Debug Mode enable bit in Debug Control register 0 (DBCR0[EDM]). Setting DBCR0[EDM] overrides the Internal Debug Mode enable bit DBCR0[IDM]. When the Hardware Debug facility is enabled, software is blocked from modifying the debug facilities. In addition, because resources are “owned” by the Hardware debugger, inconsistent values may be present if software attempts to read debug-related resources.

When hardware debug is enabled by setting DBCR0[EDM]=1, the registers and resources described in [Section 8.3, “Debug Registers,”](#) are reserved for use by the external debugger. The same events described in [Section 8.2, “Software Debug Events and Exceptions,”](#) are also used for external debugging, but exceptions are not generated to running software. Debug events enabled in the respective DBCR[0–2] registers are recorded in the DBSR regardless of MSR[DE], and no debug interrupts are generated. Instead, the CPU enters debug mode when an enabled event causes a DBSR bit to become set. DBCR0[EDM] may only be written through the OnCE port.

Access to most debug resources (registers) requires that the core clock (*m_clk*) be running in order to perform write accesses from the external hardware debugger.

Figure 8-1 shows the e200 debug resources.




Figure 8-1. e200 Debug Resources

8.2 Software Debug Events and Exceptions

Software debug events and exceptions are available when internal debug mode is enabled (DBCR0[IDM]=1) and not overridden by external debug mode (DBCR0[EDM] must be cleared). When enabled, debug events cause debug exceptions to be recorded in the Debug Status Register. Specific event types are enabled by the Debug Control Registers (DBCR0–2). The Unconditional Debug Event (UDE) is an exception to this rule; it is always enabled. Once a Debug Status Register (DBSR) bit is set (other than MRR), if Debug interrupts are enabled by MSR[DE], a Debug interrupt is generated. The debug interrupt handler is responsible for ensuring that multiple repeated debug interrupts do not occur by clearing the DBSR as appropriate.

Certain debug events are not allowed to occur when MSR[DE]=0 and DBCR0[EDM]=0. In such situations, no debug exception occurs and thus no DBSR bit is set. Other debug events may cause debug exceptions and set DBSR bits regardless of the state of MSR[DE]. A Debug interrupt is delayed until MSR[DE] is later set to '1'.

When a Debug Status Register bit is set while MSR[DE]=0 and DBCR0[EDM]=0, an Imprecise Debug Event flag (DBSR[IDE]) also is set to indicate that an exception bit in the Debug Status Register was set while Debug interrupts were disabled. Debug interrupt handler software can use this bit to determine whether the address recorded in Debug Save/Restore Register 0 is an address associated with the instruction causing the debug exception, or the address of the instruction which enabled a delayed Debug interrupt by setting the MSR[DE] bit. A **mtmsr** or **mtdbcr0** which causes both MSR[DE] and DBCR0[IDM] to become set, enabling precise debug mode, may cause an Imprecise (Delayed) Debug exception to be generated due to an earlier recorded event in the Debug Status register.

There are eight types of debug events defined by Power Architecture Book E, as follows:

1. Instruction Address Compare debug events
2. Data Address Compare debug events
3. Trap debug events
4. Branch Taken debug events
5. Instruction Complete debug events
6. Interrupt Taken debug events
7. Return debug events
8. Unconditional debug events

In addition, e200 defines additional debug events:

- The External debug events DEVT1 and DEVT2 which are described in [Section 8.2.11, “External Debug Event.”](#)
- The Critical Interrupt Taken debug event CIRPT which is described in [Section 8.2.8, “Critical Interrupt Taken Debug Event.”](#)
- The Critical Return debug event CRET which is described in [Section 8.2.10, “Critical Return Debug Event.”](#)

The e200 debug configuration supports most of these event types. Unsupported Power Architecture Book E defined functionality is as follows:

- Instruction Address Compare and Data Address Compare real address mode is not supported.
- Data Value Compare Mode is not supported.

A brief description of each of the event types follows. In these descriptions, DSRR0 and DSRR1 are used, assuming that the Debug APU is enabled. If it is disabled, use CSRR0 and CSRR1, respectively.

8.2.1 Instruction Address Compare Event

Instruction Address Compare debug events occur when enabled and execution is attempted of an instruction at an address that meets the criteria specified in the DBCR0, DBCR1, IAC1, IAC2, IAC3, and IAC4 Registers. Instruction Address compares may specify user/supervisor mode and instruction space

(MSR[IS]), along with an effective address, masked effective address, or range of effective addresses for comparison. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. IAC events do not occur when an instruction would not have normally begun execution due to a higher priority exception at an instruction boundary.

IAC compares perform a 31-bit compare for VLE instructions. Each halfword fetched by the instruction fetch unit is marked with a set of bits indicating whether an Instruction Address Compare occurred on that halfword. Debug exceptions occur if enabled and a 16-bit instruction, or the first halfword of a 32-bit instruction, is tagged with an IAC hit.

8.2.2 Data Address Compare Event

Data Address Compare debug events occur when enabled and execution of a load or store class instruction results in a data access that meets the criteria specified in the DBCR0, DBCR2, DAC1, and DAC2 Registers. Data address compares may specify user/supervisor mode and data space (MSR[DS]), along with an effective address, masked effective address, or range of effective addresses for comparison. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. Two address compare values (DAC1, DAC2) are provided.

NOTE

In contrast to the Power Architecture Book E definition, Data Address Compare events on e200 do not prevent the load or store class instruction from completing. If a load or store class instruction completes successfully without a Data TLB or Data Storage interrupt, Data Address Compare exceptions are reported at the completion of the instruction. If the exception results in a precise Debug interrupt, the address value saved in DSRR0 (or CSRR0 if the Debug APU is disabled) is the address of the instruction following the load or store class instruction.

If a load or store class instruction does not complete successfully due to a Data Storage exception, and a Data Address Compare debug exception also occurs, the result is an imprecise Debug interrupt, the address value saved in DSRR0 (or CSRR0 if the Debug APU is disabled) is the address of the load or store class instruction, and the DBSR[IDE] bit is set. In addition to occurring when DBCR0[IDM]=1, this circumstance can also occur when DBCR0[EDM]=1.

NOTE

DAC events are not recorded or counted if a **lmw** or **stmw** instruction is interrupted prior to completion by a critical input or external input interrupt.

NOTE

DAC events are not signaled on the second portion of a misaligned load or store that is broken up into two separate accesses.

8.2.3 Linked Instruction Address and Data Address Compare Event

Data Address Compare debug events may be ‘linked’ with an Instruction Address Compare event by setting the DAC1LNK and/or DAC2LNK control bits in DBCR2 to further refine when a Data Address Compare debug event is generated. DAC1 may be linked with IAC1, and DAC2 (when not used as a mask or range bounds register) may be linked with IAC3. When linked, a DAC1 (or DAC2) debug event occurs when the same instruction which generates the DAC1 (or DAC2) ‘hit’ also generates an IAC1 (or IAC3) ‘hit’. When linked, the IAC1 (or IAC3) event is not recorded in the Debug Status register, regardless of whether a corresponding DAC1 (or DAC2) event occurs, or whether the IAC1 (or IAC3) event enable is set.

When enabled and execution of a load or store class instruction results in a data access with an address that meets the criteria specified in the DBCR0, DBCR2, DAC1, and DAC2 Registers, and the instruction also meets the criteria for generating an Instruction Address Compare event, a Linked Data Address Compare debug event occurs. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. The normal DAC1 and DAC2 status bits in the DBSR are used for recording these events. The IAC1 and IAC3 status bits are not set if the corresponding Instruction Address Compare register is linked.

Linking is enabled using control bits in DBCR2.

NOTE

Linked DAC events are not recorded if a load multiple word or store multiple word instruction is interrupted prior to completion by a critical input or external input interrupt.

8.2.4 Trap Debug Event

A Trap debug event (TRAP) occurs if Trap debug events are enabled (DBCR0[TRAP]=1), a Trap instruction (**tw**) is executed, and the conditions specified by the instruction for the trap are met. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. When a Trap debug event occurs, the DBSR[TRAP] bit is set to 1 to record the debug exception.

8.2.5 Branch Taken Debug Event

A Branch Taken debug event (BRT) occurs if Branch Taken debug events are enabled (DBCR0[BRT]=1) and execution is attempted of a branch instruction, which is taken (either an unconditional branch, or a conditional branch whose branch condition is true), and MSR[DE]=1 or DBCR0[EDM]=1. Branch Taken debug events are not recognized if MSR[DE]=0 and DBCR0[EDM]=0 at the time of execution of the branch instruction and thus DBSR[IDE] can not be set by a Branch Taken debug event. When a Branch Taken debug event is recognized, the DBSR[BRT] bit is set to 1 to record the debug exception, and the address of the branch instruction is recorded in DSRR0.

8.2.6 Instruction Complete Debug Event

An Instruction Complete debug event (ICMP) occurs if Instruction Complete debug events are enabled (DBCR0[ICMP]=1), execution of any instruction is completed, and MSR[DE]=1 or DBCR0[EDM]=1. If execution of an instruction is suppressed due to the instruction causing some other exception which is

enabled to generate an interrupt, then the attempted execution of that instruction does not cause an Instruction Complete debug event. The **sc** instruction does not fall into the category of an instruction whose execution is suppressed, because the instruction actually executes and then generates a System Call interrupt. In this case, the Instruction Complete debug exception is also set. When an Instruction Complete debug event is recognized, DBSR[ICMP] is set to 1 to record the debug exception and the address of the next instruction to be executed is recorded in DSRR0.

Instruction Complete debug events are not recognized if MSR[DE]=0 and DBCR0[EDM]=0 at the time of execution of the instruction, thus DBSR[IDE] is not generally set by an ICMP debug event.

NOTE

Instruction complete debug events are not generated by the execution of an instruction which sets MSR[DE] to '1' while DBCR0[ICMP]=1, nor by the execution of an instruction which sets DBCR0[ICMP] to '1' while MSR[DE]=1 or DBCR0[EDM]=1.

8.2.7 Interrupt Taken Debug Event

An Interrupt Taken debug event (IRPT) occurs if Interrupt Taken debug events are enabled (DBCR0[IRPT]=1) and a non-critical interrupt occurs. Only non-critical class interrupts cause an Interrupt Taken debug event. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. When an Interrupt Taken debug event occurs, the DBSR[IRPT] bit is set to 1 to record the debug exception. The value saved in DSRR0 is the address of the non-critical interrupt handler.

8.2.8 Critical Interrupt Taken Debug Event

A Critical Interrupt Taken debug event (CIRPT) occurs if Critical Interrupt Taken debug events are enabled (DBCR0[CIRPT]=1) and a critical interrupt (other than a Debug interrupt when the Debug APU is disabled) occurs. Only critical class interrupts cause a Critical Interrupt Taken debug event. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. When a Critical Interrupt Taken debug event occurs, the DBSR[CIRPT] bit is set to 1 to record the debug exception. The value saved in DSRR0 is the address of the critical interrupt handler. Note that this debug event should not normally be enabled unless the Debug APU is also enabled to avoid corruption of CSRR0/1.

8.2.9 Return Debug Event

A Return debug event (RET) occurs if Return debug events are enabled (DBCR0[RET]=1) and an attempt is made to execute an **se_rfi** instruction. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. When a Return debug event occurs, the DBSR[RET] bit is set to 1 to record the debug exception.

If MSR[DE]=0 and DBCR0[EDM]=0 at the time of the execution of the **se_rfi** (such as before the MSR is updated by the **se_rfi**), then DBSR[IDE] is also set to 1 to record the imprecise debug event.

If MSR[DE]=1 at the time of the execution of the **se_rfi**, a Debug interrupt occurs provided there exists no higher priority exception which is enabled to cause an interrupt. Debug Save/Restore Register 0 is set to the address of the **se_rfi** instruction.

8.2.10 Critical Return Debug Event

A Critical Return debug event (CRET) occurs if Critical Return debug events are enabled (DBCR0[CRET]=1) and an attempt is made to execute an **se_rfc**i instruction. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. When a Critical Return debug event occurs, the DBSR[CRET] bit is set to 1 to record the debug exception.

If MSR[DE]=0 and DBCR0[EDM]=0 at the time of the execution of the **se_rfc**i (such as before the MSR is updated by the **se_rfc**i), then DBSR[IDE] is also set to 1 to record the imprecise debug event.

If MSR[DE]=1 at the time of the execution of the **se_rfc**i, a Debug interrupt occurs provided there exists no higher priority exception which is enabled to cause an interrupt. Debug Save/Restore Register 0 is set to the address of the **se_rfc**i instruction. Note that this debug event should not normally be enabled unless the Debug APU is also enabled to avoid corruption of CSRR0/1.

8.2.11 External Debug Event

An External debug event (DEVT1, DEVT2) occurs if External debug events are enabled (DBCR0[DEVT1]=1 or DBCR0[DEVT2]=1), and the respective *p_devt1* or *p_devt2* input signal transitions to the asserted state. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. When an External debug event occurs, DBSR[DEVT{1,2}] is set to '1' to record the debug exception.

8.2.12 Unconditional Debug Event

An Unconditional debug event (UDE) occurs when the Unconditional Debug Event (*p_ude*) input transitions to the asserted state, and either DBCR0[IDM]=1 or DBCR0[EDM]=1. The Unconditional debug event is the only debug event which does not have a corresponding enable bit for the event in DBCR0. This event can occur and be recorded in DBSR regardless of the setting of MSR[DE]. When an Unconditional debug event occurs, the DBSR[UDE] bit is set to '1' to record the debug exception.

8.3 Debug Registers

This section describes debug-related registers that are software accessible. These registers are intended for use by special debug tools and debug software, not by general application code.

Access to these registers by software is conditioned by the External Debug Mode control bit (DBCR0[EDM]) which can be set by the hardware debug port. If DBCR0[EDM] is set, software is prevented from modifying debug register values. Execution of a **mtspr** instruction targeting a debug register does not cause modifications to occur. In addition, because the external debugger hardware may be manipulating debug register values, the state of these registers is not guaranteed to be consistent if accessed (read) by software with a **mfspr** instruction, other than the DBCR0[EDM] bit itself.

8.3.1 Debug Address and Value Registers

Instruction Address Compare registers IAC1, IAC2, IAC3, and IAC4 are used to hold instruction addresses for address comparison purposes. In addition, IAC2 and IAC4 hold mask information for IAC1

and IAC3 respectively when *Address Bit Match* compare modes are selected. Note that when performing instruction address compares, the low order bit of the instruction address and the corresponding IAC register is ignored.

Data Address Compare registers DAC1 and DAC2 are used to hold data access addresses for address comparison purposes. In addition, DAC2 holds mask information for DAC1 when *Address Bit Match* compare mode is selected.

8.3.2 Debug Control and Status Registers

Debug Control Registers 0-2 (DBCR0, DBCR1, DBCR2) are used to enable debug events, reset the processor, and set the debug mode of the processor. The Debug Status register (DBSR) records debug exceptions while Internal or External Debug Mode is enabled.

e200 requires that a context synchronizing instruction follow a **mtspr** DBCR0-2 or DBSR to ensure that any alterations enabling/disabling debug events are effective. The context synchronizing instruction may or may not be affected by the alteration. Typically, an **isync** instruction is used to create a synchronization boundary beyond which it can be guaranteed that the newly written control values are in effect.

For watchpoint generation, configuration settings contained in DBCR1 and DBCR2 are used, even though the corresponding event(s) may be disabled (via DBCR0) from setting DBSR flags.

8.3.2.1 Debug Control Register 0 (DBCR0)

Debug Control Register 0 is used to enable debug modes and controls which debug events are allowed to set DBSR flags. e200 adds some implementation specific bits to this register, as seen in [Figure 8-2](#).

EDM	IDM	RST		ICMP	BRT	IRPT	TRAP	IAC1	IAC2	IAC3	IAC4	DAC1	DAC2	RET	0					DEVT1	DEVT2	0			CIRPT	CRET	0				
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

SPR—308: Read/Write; Reset¹—0x0

Figure 8-2. DBCR0 Register

¹ DBCR0[EDM] is affected by *j_trst_b* or *m_por* assertion, and while in the Test_Logic_Reset state, but not by *p_reset_b*. All other bits are reset by processor reset *p_reset_b* as well as by *m_por*.

Table 8-1 provides bit definitions for Debug Control Register 0.

Table 8-1. DBCR0 Bit Definitions

Bit(s)	Name	Description
0	EDM	External Debug Mode. This bit is read-only by software. 0 External debug mode disabled. Internal debug events not mapped into external debug events. 1 External debug mode enabled. Events do not cause the CPU to vector to interrupt code. Software is not permitted to write to debug registers {DBCRx, DBSR, DBCNT, IAC1-4, DAC1-2}. Programming Notes: It is recommended that debug status bits in the Debug Status Register be cleared before disabling external debug mode to avoid any internal imprecise debug interrupts. Software may use this bit to determine if external debug has control over the debug registers. The hardware debugger must set the EDM bit to '1' before other bits in this register (and other debug registers) may be altered. On the initial setting of this bit to '1', all other bits are unchanged. This bit is only writable through the OnCE port.
1	IDM	Internal Debug Mode 0 Debug exceptions are disabled. Debug events do not affect DBSR unless EDM is set. 1 Debug exceptions are enabled. Enabled debug events update the DBSR. If MSR[DE]=1, the occurrence of a debug event, or the recording of an earlier debug event in the Debug Status Register when MSR[DE] was cleared, causes a Debug interrupt.
2:3	RST	Reset Control 00 No function 01 Reserved 10 <i>p_resetout_b</i> pin asserted by Debug Reset Control. Allows external device to initiate processor reset. 11 Reserved
4	ICMP	Instruction Complete Debug Event Enable 0 ICMP debug events are disabled 1 ICMP debug events are enabled
5	BRT	Branch Taken Debug Event Enable 0 BRT debug events are disabled 1 BRT debug events are enabled
6	IRPT	Interrupt Taken Debug Event Enable 0 IRPT debug events are disabled 1 IRPT debug events are enabled
7	TRAP	Trap Taken Debug Event Enable 0 TRAP debug events are disabled 1 TRAP debug events are enabled
8	IAC1	Instruction Address Compare 1 Debug Event Enable 0 IAC1 debug events are disabled 1 IAC1 debug events are enabled
9	IAC2	Instruction Address Compare 2 Debug Event Enable 0 IAC2 debug events are disabled 1 IAC2 debug events are enabled
10	IAC3	Instruction Address Compare 3 Debug Event Enable 0 IAC3 debug events are disabled 1 IAC3 debug events are enabled

Table 8-1. DBCR0 Bit Definitions (continued)

Bit(s)	Name	Description
11	IAC4	Instruction Address Compare 4 Debug Event Enable 0 IAC4 debug events are disabled 1 IAC4 debug events are enabled
12:13	DAC1	Data Address Compare 1 Debug Event Enable 00 DAC1 debug events are disabled 01 DAC1 debug events are enabled only for store-type data storage accesses 10 DAC1 debug events are enabled only for load-type data storage accesses 11 DAC1 debug events are enabled for load-type or store-type data storage accesses
14:15	DAC2	Data Address Compare 2 Debug Event Enable 00 DAC2 debug events are disabled 01 DAC2 debug events are enabled only for store-type data storage accesses 10 DAC2 debug events are enabled only for load-type data storage accesses 11 DAC2 debug events are enabled for load-type or store-type data storage accesses
16	RET	Return Debug Event Enable 0 RET debug events are disabled 1 RET debug events are enabled
17:20	—	Reserved
21	DEVT1	External Debug Event 1 Enable 0 DEVT1 debug events are disabled 1 DEVT1 debug events are enabled
22	DEVT2	External Debug Event 2 Enable 0 DEVT2 debug events are disabled 1 DEVT2 debug events are enabled
23:24	—	Reserved
25	CIRPT	Critical Interrupt Taken Debug Event Enable 0 CIRPT debug events are disabled 1 CIRPT debug events are enabled
26	CRET	Critical Return Debug Event Enable 0 CRET debug events are disabled 1 CRET debug events are enabled
27:31	—	Reserved

8.3.2.2 Debug Control Register 1 (DBCR1)

Debug Control Register 1 is used to configure Instruction Address Compare operation. The DBCR1 register is shown in [Figure 8-3](#).



Figure 8-3. DBCR1 Register

Table 8-2 provides bit definitions for Debug Control Register 1.

Table 8-2. DBCR1 Bit Definitions

Bit(s)	Name	Description
0:1	IAC1US	Instruction Address Compare 1 User/Supervisor Mode 00 IAC1 debug events not affected by MSR[PR] 01 Reserved 10 IAC1 debug events can only occur if MSR[PR]=0 (Supervisor mode) 11 IAC1 debug events can only occur if MSR[PR]=1. (User mode)
2:3	IAC1ER	Instruction Address Compare 1 Effective/Real Mode 00 IAC1 debug events are based on effective address 01 Unimplemented in e200 (Book E real address compare), no match can occur 10 IAC1 debug events are based on effective address and can only occur if MSR[IS]=0 11 IAC1 debug events are based on effective address and can only occur if MSR[IS]=1
4:5	IAC2US	Instruction Address Compare 2 User/Supervisor Mode 00 IAC2 debug events not affected by MSR[PR] 01 Reserved 10 IAC2 debug events can only occur if MSR[PR]=0 (Supervisor mode) 11 IAC2 debug events can only occur if MSR[PR]=1. (User mode)
6:7	IAC2ER	Instruction Address Compare 2 Effective/Real Mode 00 IAC2 debug events are based on effective address 01 Unimplemented in e200 (Book E real address compare), no match can occur 10 IAC2 debug events are based on effective address and can only occur if MSR[IS]=0 11 IAC2 debug events are based on effective address and can only occur if MSR[IS]=1
8:9	IAC12M	Instruction Address Compare 1/2 Mode 00 Exact address compare. IAC1 debug events can only occur if the address of the instruction fetch is equal to the value specified in IAC1. IAC2 debug events can only occur if the address of the instruction fetch is equal to the value specified in IAC2. 01 Address bit match. IAC1 debug events can occur only if the address of the instruction fetch, ANDed with the contents of IAC2 are equal to the contents of IAC1, also ANDed with the contents of IAC2. IAC2 debug events do not occur. IAC1US and IAC1ER settings are used. 10 Inclusive address range compare. IAC1 debug events can occur only if the address of the instruction fetch is greater than or equal to the value specified in IAC1 and less than the value specified in IAC2. IAC2 debug events do not occur. IAC1US and IAC1ER settings are used. 11 Exclusive address range compare. IAC1 debug events can occur only if the address of the instruction fetch is less than the value specified in IAC1 or is greater than or equal to the value specified in IAC2. IAC2 debug events do not occur. IAC1US and IAC1ER settings are used.
10:15	—	Reserved
16:17	IAC3US	Instruction Address Compare 3 User/Supervisor Mode 00 IAC3 debug events not affected by MSR[PR] 01 Reserved 10 IAC3 debug events can only occur if MSR[PR]=0 (Supervisor mode) 11 IAC3 debug events can only occur if MSR[PR]=1 (User mode)
18:19	IAC3ER	Instruction Address Compare 3 Effective/Real Mode 00 IAC3 debug events are based on effective address 01 Unimplemented in e200 (Book E real address compare), no match can occur 10 IAC3 debug events are based on effective address and can only occur if MSR[IS]=0 11 IAC3 debug events are based on effective address and can only occur if MSR[IS]=1

Table 8-2. DBCR1 Bit Definitions (continued)

Bit(s)	Name	Description
20:21	IAC4US	Instruction Address Compare 4 User/Supervisor Mode 00 IAC4 debug events not affected by MSR[PR] 01 Reserved 10 IAC4 debug events can only occur if MSR[PR]=0 (Supervisor mode). 11 IAC4 debug events can only occur if MSR[PR]=1. (User mode)
22:23	IAC4ER	Instruction Address Compare 4Effective/Real Mode 00 IAC4 debug events are based on effective address 01 Unimplemented in e200 (Book E real address compare), no match can occur 10 IAC4 debug events are based on effective address and can only occur if MSR[IS]=0 11 IAC4 debug events are based on effective address and can only occur if MSR[IS]=1
24:25	IAC34M	Instruction Address Compare 3/4 Mode 00 Exact address compare. IAC3 debug events can only occur if the address of the instruction fetch is equal to the value specified in IAC3. IAC4 debug events can only occur if the address of the instruction fetch is equal to the value specified in IAC4. 01 Address bit match. IAC3 debug events can occur only if the address of the instruction fetch, ANDed with the contents of IAC4 are equal to the contents of IAC3, also ANDed with the contents of IAC4. IAC4 debug events do not occur. IAC3US and IAC3ER settings are used. 10 Inclusive address range compare. IAC3 debug events can occur only if the address of the instruction fetch is greater than or equal to the value specified in IAC3 and less than the value specified in IAC4. IAC4 debug events do not occur. IAC3US and IAC3ER settings are used. 11 Exclusive address range compare. IAC3 debug events can occur only if the address of the instruction fetch is less than the value specified in IAC3 or is greater than or equal to the value specified in IAC4. IAC4 debug events do not occur. IAC3US and IAC3ER settings are used.
26:31	—	Reserved

8.3.2.3 Debug Control Register 2 (DBCR2)

Debug Control Register 2 is used to configure Data Address Compare and Data Value Compare operation. The DBCR2 register is shown in [Figure 8-4](#).




Figure 8-4. DBCR2 Register

Table 8-3 provides bit definitions for Debug Control Register 2.

Table 8-3. DBCR2 Bit Definitions

Bit(s)	Name	Description
0:1	DAC1US	Data Address Compare 1 User/Supervisor Mode 00 DAC1 debug events not affected by MSR[PR] 01 Reserved 10 DAC1 debug events can only occur if MSR[PR]=0 (Supervisor mode) 11 DAC1 debug events can only occur if MSR[PR]=1. (User mode)
2:3	DAC1ER	Data Address Compare 1 Effective/Real Mode 00 DAC1 debug events are based on effective address 01 Unimplemented in e200 (Book E real address compare), no match can occur 10 DAC1 debug events are based on effective address and can only occur if MSR[DS]=0 11 DAC1 debug events are based on effective address and can only occur if MSR[DS]=1
4:5	DAC2US	Data Address Compare 2 User/Supervisor Mode. 00 DAC2 debug events not affected by MSR[PR] 01 Reserved 10 DAC2 debug events can only occur if MSR[PR]=0 (Supervisor mode) 11 DAC2 debug events can only occur if MSR[PR]=1. (User mode)
6:7	DAC2ER	Data Address Compare 2 Effective/Real Mode 00 DAC2 debug events are based on effective address 01 Unimplemented in e200 (Book E real address compare), no match can occur 10 DAC2 debug events are based on effective address and can only occur if MSR[DS]=0 11 DAC2 debug events are based on effective address and can only occur if MSR[DS]=1
8:9	DAC12M	Data Address Compare 1/2 Mode 00 Exact address compare. DAC1 debug events can only occur if the address of the data access is equal to the value specified in DAC1. DAC2 debug events can only occur if the address of the data access is equal to the value specified in DAC2. 01 Address bit match. DAC1 debug events can occur only if the address of the data access ANDed with the contents of DAC2, are equal to the contents of DAC1 also ANDed with the contents of DAC2. DAC2 debug events do not occur. DAC1US and DAC1ER settings are used. 10 Inclusive address range compare. DAC1 debug events can occur only if the address of the data access is greater than or equal to the value specified in DAC1 and less than the value specified in DAC2. DAC2 debug events do not occur. DAC1US and DAC1ER settings are used. 11 Exclusive address range compare. DAC1 debug events can occur only if the address of the data access is less than the value specified in DAC1 or is greater than or equal to the value specified in DAC2. DAC2 debug events do not occur. DAC1US and DAC1ER settings are used.
10	DAC1LNK	Data Address Compare 1 Linked 0 no affect 1 DAC1 debug events are linked to IAC1 debug events. IAC1 debug events do not affect DBSR When linked to IAC1, DAC1 debug events are conditioned based on whether the instruction also generated an IAC1 debug event
11	DAC2LNK	Data Address Compare 2 Linked 0 no affect 1 DAC 2 debug events are linked to IAC3 debug events. IAC3 debug events do not affect DBSR When linked to IAC3, DAC2 debug events are conditioned based on whether the instruction also generated an IAC3 debug event. DAC2 can only be linked if DAC12M specifies <i>Exact Address Compare</i> because DAC2 debug events are not generated in the other compare modes.
12:31	—	Reserved for Data Value Compare control (not supported by e200)

8.3.2.4 Debug Status Register (DBSR)

The Debug Status Register (DBSR) contains status on debug events and the most recent processor reset. The Debug Status Register is set via hardware, and read and cleared via software. Bits in the Debug Status Register can be cleared using **mtspr DBSR,RS**. Clearing is done by writing to the Debug Status Register with a 1 in any bit position that is to be cleared and 0 in all other bit positions. The write data to the Debug Status Register is not direct data, but a mask. A ‘1’ causes the bit to be cleared, and a ‘0’ has no effect. Debug Status bits are set by Debug events only while Internal Debug Mode is enabled or External Debug Mode is enabled. When debug interrupts are enabled (MSR[DE]=1, DBCR0[IDM]=1, and DBCR0[EDM]=0), a set bit in DBSR other than MRR or VLES causes a debug interrupt to be generated. The debug interrupt handler is responsible for clearing DBSR bits prior to returning to normal execution. The Power Architecture VLE APU adds the DBSR[VLES] status bit to indicate debug events occurring due to a Power Architecture VLE instruction. The DBSR register is shown in [Figure 8-5](#).




Figure 8-5. DBSR Register

[Table 8-4](#) provides bit definitions for the Debug Status Register.

Table 8-4. DBSR Bit Definitions

Bit(s)	Name	Description
0	IDE	Imprecise Debug Event Set to 1 if MSR[DE]=0 and DBCR0[EDM]=0 and a debug event causes its respective Debug Status Register bit to be set to 1. It may also be set to ‘1’ if DBCR0[EDM]=1 and an imprecise debug event occurs due to a DAC event on a load or store which is terminated with error.
1	UDE	Unconditional Debug Event Set to 1 if an Unconditional debug event occurred.
2:3	MRR	Most Recent Reset. 00 No reset occurred because these bits were last cleared by software 01 A hard reset occurred because these bits were last cleared by software 10 Reserved 11 Reserved
4	ICMP	Instruction Complete Debug Event Set to 1 if an Instruction Complete debug event occurred.
5	BRT	Branch Taken Debug Event Set to 1 if an Branch Taken debug event occurred.
6	IRPT	Interrupt Taken Debug Event Set to 1 if an Interrupt Taken debug event occurred.
7	TRAP	Trap Taken Debug Event Set to 1 if a Trap Taken debug event occurred.

Table 8-4. DBSR Bit Definitions (continued)

Bit(s)	Name	Description
8	IAC1	Instruction Address Compare 1 Debug Event Set to 1 if an IAC1 debug event occurred.
9	IAC2	Instruction Address Compare 2 Debug Event Set to 1 if an IAC2 debug event occurred.
10	IAC3	Instruction Address Compare 3 Debug Event Set to 1 if an IAC3 debug event occurred.
11	IAC4	Instruction Address Compare 4 Debug Event Set to 1 if an IAC4 debug event occurred.
12	DAC1R	Data Address Compare 1 Read Debug Event Set to 1 if a read-type DAC1 debug event occurred while DBCR0[DAC1]=0b10 or DBCR0[DAC1]=0b11
13	DAC1W	Data Address Compare 1 Write Debug Event Set to 1 if a write-type DAC1 debug event occurred while DBCR0[DAC1]=0b01 or DBCR0[DAC1]=0b11
14	DAC2R	Data Address Compare 2 Read Debug Event Set to 1 if a read-type DAC2 debug event occurred while DBCR0[DAC2]=0b10 or DBCR0[DAC2]=0b11
15	DAC2W	Data Address Compare 2 Write Debug Event Set to 1 if a write-type DAC2 debug event occurred while DBCR0[DAC2]=0b01 or DBCR0[DAC2]=0b11
16	RET	Return Debug Event Set to 1 if a Return debug event occurred
17:20	—	Reserved
21	DEVT1	External Debug Event 1 Debug Event Set to 1 if a DEVT1 debug event occurred
22	DEVT2	External Debug Event 2 Debug Event Set to 1 if a DEVT2 debug event occurred
23:24	—	Reserved
25	CIRPT	Critical Interrupt Taken Debug Event Set to 1 if a Critical Interrupt Taken debug event occurred.
26	CRET	Critical Return Debug Event Set to 1 if a Critical Return debug event occurred
27	VLES	VLE Status Set to 1 if an ICMP, BRT, TRAP, RET, CRET, IAC, or DAC debug event occurred on a Power Architecture VLE Instruction. Undefined for IRPT, CIRPT, DEVT[1,2], and UDE events
28:31	—	Reserved

8.4 External Debug Support

External debug support is supplied through the e200 OnCE controller serial interface which allows access to internal core registers and other system state while in External Debug Mode (EDM). All debug resources including DBCR0-2, DBSR, IAC1-4, and DAC1-2 are accessible through the serial OnCE interface in external debug mode. Setting the DBCR0[EDM]bit to ‘1’ through the OnCE interface enables external debug mode and disables software updates to the debug registers. When DBCR0[EDM] is set, debug events enabled to set respective DBSR status bits also cause the CPU to enter Debug Mode as

opposed to generating Debug Interrupts. In Debug Mode, the CPU is halted at a recoverable boundary, and an external Debug Control Module may control CPU operation through the On-Chip Emulation logic (OnCE). No Debug interrupts can occur while DBCR0[EDM] remains set.

NOTE

On the initial setting of DBCR0[EDM] to '1', other bits in DBCR0 remain unchanged. After DBCR0[EDM] has been set, all debug register resources may be subsequently controlled through the OnCE interface. The DBSR register should be cleared as part of the process of enabling external debug activity. The core should be placed into debug mode via the OCR[DR] control bit prior to writing EDM to '1'. This gives the debugger the opportunity to cleanly write to the DBCRx registers and the DBSR to clear out any residual state / control information which could cause unintended operation.

NOTE

It is intended for the core to remain in external debug mode (DBCR0[EDM]=1) in order to single step or perform other debug mode entry/ reentry via the OCR[DR], by performing go+noexit commands, or by assertion of the *jd_de_b* signal.

NOTE

DBCR0[EDM] operation is blocked if OnCE operation is disabled (*jd_en_once* negated) regardless of whether it is set or cleared. This means that if DBCR0[EDM] was previously set, and then *jd_en_once* is negated (this should not occur), entry into debug mode is blocked, all events are blocked, and watchpoints are blocked.

Due to clock domain design, the CPU clock (*m_clk*) must be active in order to perform writes to debug registers other than the OnCE Command register (OCMD), the OnCE Control register (OCR), or the DBCR0[EDM] bit. Register read data is synchronized back to the *j_tclk* clock domain. The OnCE Control register provides the capability of signaling the system level clock controller that the CPU clock should be activated if not already active.

Updates to the DBCRx and DBSR registers via the OnCE interface should be performed with the CPU in debug mode to guarantee proper operation. Due to the various points in the CPU pipeline where control is sampled and event handshaking is performed, it is possible that modifications to these registers while the CPU is running may result in early or late entry into debug mode, and may have incorrect status posted in the DBSR register.

8.4.1 OnCE Introduction

The e200 on-chip emulation circuitry (OnCE™/Nexus Class 1 interface) provides a means of interacting with the e200 core and integrated system so that a user may examine registers, memory, or on-chip peripherals facilitating hardware/software development. OnCE operation is controlled via an industry standard IEEE 1149.1 TAP controller. By using public instructions, the external hardware debugger can freeze or halt the CPU, read and write internal state, and resume normal execution. The core does not

contain IEEE 1149.1 standard boundary cells on its interface, as it is a building block for further integration. It does not support the JTAG related boundary scan instruction functionality, although JTAG public instructions may be decoded and signaled to external logic.

The OnCE logic provides for Nexus Class 1 static debug capability (utilizing the same set of resources available to software while in internal debug mode), and is present in all e200-based designs. The OnCE module also provides support for directly integrating a Nexus class 2 or class 3 Real-Time Debug unit with the e200 core for development of real-time systems where traditional static debug is insufficient. The partitioning between a OnCE module and a connected Nexus module to provide real-time debug allows for capability and cost trade-offs to be made.

The e200 core is designed to be a fully integrateable module. The OnCE TAP controller and associated enabling logic are designed to allow concatenation with an existing JTAG controller if present in the system. Thus, the e200 module can be easily integrated with existing JTAG designs or as a stand-alone controller.

In order to enable full OnCE operation, the *jd_enable_once* input signal must be asserted. In some system integrations, this is automatic, because the input is tied asserted. Other integrations may require the execution of the Enable OnCE command via the TAP and appropriate entry of serial data. Exact requirements are documented by the integrated product specification. The *jd_enable_once* input signal should not change state during a debug session, or undefined activity may occur.

The following figures show the TAP controller state model and the TAP registers implemented by the OnCE logic.




Figure 8-6. OnCE TAP Controller and Registers

The OnCE controller is implemented as a 16-state FSM, with a one-to-one correspondence to the states defined for the JTAG TAP controller.

Instructions may be single-stepped by scanning new values into the CPUSCR, and performing a OnCE go+noexit command (See [Section 8.4.5.2, “e200 OnCE Command Register \(OCMD\)”](#)). The CPU then temporarily exits the debug state (but not the debug session) to execute the instruction, and then returns to the debug state (again indicated via the OnCE Status Register (OSR)). The debug session remains in force until the final OnCE go+exit command is executed, at which time the CPU returns to the previous state it was in (unless a new debug request is pending). A scan into the CPUSCR is required prior to executing each go+exit or go+noexit OnCE command.

8.4.2 JTAG/OnCE Pins

The JTAG/OnCE pin interface is used to transfer OnCE instructions and data to the OnCE control block. Depending on the particular resource being accessed, the CPU may need to be placed in the Debug mode. For resources outside of the CPU block and contained in the OnCE block, the processor is not disturbed, and may continue execution. If a processor resource is required, an internal debug request (*dbg_dbgrq*) may be asserted to the CPU by the OnCE controller, and causes the CPU to finish the current instruction being executed, save the instruction pipeline information, enter Debug Mode, and wait for further commands. Asserting *dbg_dbgrq* causes the chip to exit the low power mode enabled by the setting of MSR[WE], as well as temporarily exiting the waiting, stopped, or halted power management states.

[Table 8-5](#) details the primary JTAG/OnCE interface signals.

Table 8-5. JTAG/OnCE Primary Interface Signals

Signal Name	Type	Description
<i>j_trst_b</i>	I	JTAG test reset
<i>j_tclk</i>	I	JTAG test clock
<i>j_tms</i>	I	JTAG test mode select
<i>j_tdi</i>	I	JTAG test data input
<i>j_tdo</i>	O	Test data out to master controller or pad
<i>j_tdo_en</i> ¹	O	Enables TDO output buffer

¹ *j_tdo_en* is asserted when the TAP controller is in the shift_DR or shift_IR state.

A full description of JTAG pins is provided in [Section 6.3.15, “JTAG Support Signals.”](#)

8.4.3 OnCE Internal Interface Signals

The following paragraphs describe the e200 OnCE interface signals to other internal blocks associated with the e200 OnCE controller.

8.4.3.1 CPU Debug Request (*dbg_dbgrq*)

The *dbg_dbgrq* signal is asserted by the e200 OnCE control logic to request the CPU to enter the debug state. It may be asserted for a number of different conditions, and causes the CPU to finish the current

instruction being executed, save the instruction pipeline information, enter the debug mode, and wait for further commands.

8.4.3.2 CPU Debug Acknowledge (*cpu_dbgack*)

The *cpu_dbgack* signal is asserted by the CPU upon entering the debug state. This signal is used as part of the handshake mechanism between the e200 OnCE control logic and the rest of the CPU. The CPU core may enter debug mode either through a software or hardware event.

8.4.3.3 CPU Address, Attributes

The CPU address and attribute information are used by a Nexus class 2-4 debug unit with information for real-time address trace information.

8.4.3.4 CPU Data

The CPU data bus(es) are used to supply a Nexus class 2-4 debug unit with information for real-time data trace capability.

8.4.4 OnCE Interface Signals

The following paragraphs describe additional e200 OnCE interface signals to other external blocks such as a Nexus Controller and external blocks which may need information pertaining to debug operation.

8.4.4.1 OnCE Enable (*jd_en_once*)

The OnCE enable signal *jd_en_once* is used to enable the OnCE controller to allow certain instructions and operations to be executed. Assertion of this signal enables the full OnCE command set, as well as operation of control signals and OnCE Control register functions. When this signal is disabled, only the Bypass, ID and Enable_OnCE commands are executed by the OnCE unit, and all other commands default to a “Bypass” command. The OnCE Status register (OSR) is not visible when OnCE operation is disabled. In addition, OnCE Control register (OCR) functions are disabled, as is the operation of the *jd_de_b* input. Secure systems may choose to leave the *jd_en_once* signal negated until a security check has been performed. Other systems should tie this signal asserted to enable full OnCE operation. The *j_en_once_regssel* output signal is provided to assist external logic performing security checks. Refer to [Section 6.3.15.15, “Enable Once Register Select \(*j_en_once_regssel*\),”](#) for a description of the *j_en_once_regssel* output signal.

The *jd_en_once* input must only change state during the Test-Logic-Reset, Run-Test/Idle, or Update_DR TAP states. A new value takes effect after one additional *j_tclk* cycle of synchronization. In addition, *jd_enable_once* input signal must not change state during a debug session, or undefined activity may occur.

8.4.4.2 OnCE Debug Request/Event (*jd_de_b*, *jd_de_en*)

If implemented at the SoC level, a system level bidirectional open drain debug event pin *DE_b* (not part of the e200 interface) provides a fast means of entering the Debug Mode of operation from an external

command controller (when input) as well as a fast means of acknowledging the entering of the Debug Mode of operation to an external command controller (when output). The assertion of this pin by a command controller causes the CPU core to finish the current instruction being executed, save the instruction pipeline information, enter Debug Mode, and wait for commands to be entered. If *DE_b* was used to enter the Debug Mode then *DE_b* must be negated after the OnCE controller responds with an acknowledge and before sending the first OnCE command. The assertion of this pin by the CPU Core acknowledges that it has entered the Debug Mode and is waiting for commands to be entered.

To support operation of this system pin, the OnCE logic supplies the *jd_de_en* output and samples the *jd_de_b* input when OnCE is enabled (*jd_en_once* asserted). Assertion of *jd_de_b* causes the OnCE logic to place the CPU into Debug Mode. Once Debug Mode has been entered, the *jd_de_en* output is asserted for three *j_tclk* periods to signal an acknowledge. *jd_de_en* can be used to enable the open-drain pulldown of the system level *DE_b* pin.

For systems which do not implement a system level bidirectional open drain debug event pin *DE_b*, the *jd_de_en* and *jd_de_b* signals may still be used to handshake debug entry.

8.4.4.3 e200 OnCE Debug Output (*jd_debug_b*)

The e200 OnCE Debug output *jd_debug_b* is used to indicate to on-chip resources that a debug session is in progress. Peripherals and other units may use this signal to modify normal operation for the duration of a debug session, which may involve the CPU executing a sequence of instructions solely for the purpose of visibility/system control which are not part of the normal instruction stream the CPU would have executed had it not been placed in debug mode. This signal is asserted the first time the CPU enters the debug state, and remains asserted until the CPU is released by a write to the e200 OnCE Command Register with the GO and EX bits set, and a register specified as either “No Register Selected” or the CPUSCR. This signal remains asserted even though the CPU may enter and exit the debug state for each instruction executed under control of the e200 OnCE controller. See [Section 8.4.5.2, “e200 OnCE Command Register \(OCMD\),”](#) for more information on the function of the GO and EX bits. This signal is not normally used by the CPU.

8.4.4.4 e200 CPU Clock On Input (*jd_mclk_on*)

The e200 CPU Clock On input *jd_mclk_on* is used to indicate that the CPU’s *m_clk* input is active. This input signal is expected to be driven by system logic external to the e200 core, is synchronized to the *j_tclk* (scan clock) clock domain, and is presented as a status flag on the *j_tdo* output during the Shift_IR state. External firmware may use this signal to ensure proper scan sequences occur to access debug resources in the *m_clk* clock domain.

8.4.4.5 Watchpoint Events (*jd_watchpt[0:5]*)

The *jd_watchpt[0:5]* signals may be asserted by the e200 OnCE control logic to signal that a watchpoint condition has occurred. Watchpoints do not cause the CPU to be affected. They are provided to allow external visibility only. Watchpoint events are conditioned by the settings in the DBCR0, DBCR1, and DBCR2 registers.

8.4.5 e200 OnCE Controller and Serial Interface

The e200 OnCE Controller contains the e200 OnCE command register, the e200 OnCE decoder, and the status/control register. Figure 8-7 is a block diagram of the e200 OnCE controller. In operation, the e200 OnCE Command register acts as the IR for the e200 TAP controller, and all other OnCE resources are treated as data registers (DR) by the TAP controller. The Command register is loaded by serially shifting in commands during the TAP controller Shift-IR state, and is loaded during the Update-IR state. The Command register selects a resource to be accessed as a data register (DR) during the TAP controller Capture-DR, Shift-DR and Update-DR states.




Figure 8-7. e200 OnCE Controller and Serial Interface

8.4.5.1 e200 OnCE Status Register

Status information regarding the state of the e200 CPU is latched into the OnCE Status register when the OnCE controller state machine enters the Capture-IR state. When OnCE operation is enabled, this information is provided on the *j_tdo* output in serial fashion when the Shift_IR state is entered following a Capture-IR. Information is shifted out least significant bit first.

MCLK	ERR	CHKSTOP	RESET	HALT	STOP	DEBUG	WAIT	0	1
0	1	2	3	4	5	6	7	8	9

Figure 8-8. OnCE Status Register

Table 8-6 provides bit definitions for the Once Status Register.

Table 8-6. OnCE Status Register Bit Definitions

Bit(s)	Name	Description
0	MCLK	MCLK <i>m_clk</i> Status Bit 0 Inactive state 1 Active state This status bit reflects the logic level on the <i>jd_mclk_on</i> input signal after capture by <i>j_tclk</i> .
1	ERR	ERROR This bit is used to indicate that an error condition occurred during attempted execution of the last single-stepped instruction (GO+NoExit with CPUSCR or No Register Selected in OCMD), and that the instruction may not have been properly executed. This could occur if an Interrupt (all classes including External, Critical, machine check, Storage, Alignment, Program, etc.) occurred while attempting to perform the instruction single step. In this case, the CPUSCR contains information related to the first instruction of the Interrupt handler, and no portion of the handler will have been executed.
2	CHKSTOP	CHECKSTOP Mode This bit reflects the logic level on the CPU <i>p_chkstop</i> output after capture by <i>j_tclk</i> .
3	RESET	RESET Mode This bit reflects the <u>inverted</u> logic level on the CPU <i>p_reset_b</i> input after capture by <i>j_tclk</i> .
4	HALT	HALT Mode This bit reflects the logic level on the CPU <i>p_halted</i> output after capture by <i>j_tclk</i> .
5	STOP	STOP Mode This bit reflects the logic level on the CPU <i>p_stopped</i> output after capture by <i>j_tclk</i> .
6	DEBUG	Debug Mode This bit is asserted once the CPU is in debug mode. It is negated once the CPU exits debug mode (even during a debug session)
7	WAIT	Waiting Mode This bit reflects the logic level on the CPU <i>p_waiting</i> output after capture by <i>j_tclk</i> .
8	—	Reserved, set to 0 for 1149.1 compliance
9	—	Reserved, set to 1 for 1149.1 compliance

8.4.5.2 e200 OnCE Command Register (OCMD)

The OnCE Command Register (OCMD) is a 10-bit shift register that receives its serial data from the TDI pin and serves as the instruction register (IR). It holds the 10-bit commands to be used as input for the e200 OnCE Decoder. The Command Register is shown in Figure 8-9. The OCMD is updated when the TAP controller enters the Update-IR state. It contains fields for controlling access to a resource, as well as controlling single-step operation and exit from OnCE mode.

Although the OCMD is updated during the Update-IR TAP controller state, the corresponding resource is accessed in the DR scan sequence of the TAP controller, and as such, the Update-DR state must be transitioned through in order for an access to occur. In addition, the Update-DR state must also be transitioned through in order for the single-step and/or exit functionality to be performed, even though the command appears to have no data resource requirement associated with it.

R/W	GO	EX	RS[0:6]						
0	1	2	3	4	5	6	7	8	9

Reset—10'b1000000010 on assertion of *j_trst_b* or *m_por*, or while in the Test_Logics_Reset state

Figure 8-9. OnCE Command Register

Table 8-7 provides bit definitions for the Once Command Register.

Table 8-7. OnCE Command Register Bit Definitions

Bit(s)	Name	Description
0	R/W	Read/Write Command Bit The R/W bit specifies the direction of data transfer. The table below describes the options defined by the R/W bit. 0 Write the data associated with the command into the register specified by RS[0:6] 1 Read the data contained in the register specified by RS[0:6] Note: The R/W bit generally ignored for read-only or write-only registers. In addition, it is ignored for all bypass operations. When performing writes, most registers are sampled in the Capture-DR state into a 32-bit shift register, and subsequently shifted out on <i>j_tdo</i> during the first 32 clocks of Shift-DR.
1	GO	Go Go Command Bit 0 Inactive (no action taken) 1 Execute instruction in IR If the GO bit is set, the chip executes the instruction that resides in the IR register in the CPUSCR. To execute the instruction, the processor leaves the debug mode, executes the instruction, and if the EX bit is cleared, returns to the debug mode immediately after executing the instruction. The processor goes on to normal operation if the EX bit is set, and no other debug request source is asserted. The GO command is executed only if the operation is a read/write to CPUSCR or a read/write to “No Register Selected”. Otherwise the GO bit is ignored. The processor leaves the debug mode after the TAP controller Update-DR state is entered. On a GO+NoExit operation, returning to debug mode is treated as a debug event, thus exceptions such as machine checks and interrupts may take priority and prevent execution of the intended instruction. Debug firmware should mask these exceptions as appropriate. The OSR[ERR] bit indicates such an occurrence.
2	EX	Exit Command Bit 0 Remain in debug mode 1 Leave debug mode If the EX bit is set, the processor leaves the debug mode and resume normal operation until another debug request is generated. The Exit command is executed only if the Go command is issued, and the operation is a read/write to CPUSCR or a read/write to “No Register Selected”. Otherwise the EX bit is ignored. The processor leaves the debug mode after the TAP controller Update-DR state is entered. Note that if the DR bit in the OnCE control register is set or remains set, or if a bit in the DBSR is set and DBCR0[EDM]=1 (external debug mode is enabled), or if another debug request source is asserted, then the processor may return to the debug mode <u>without</u> execution of an instruction, even though the EX bit was set.
3:9	RS	Register Select The Register Select bits define which register is source (destination) for the read (write) operation. Table 8-9 indicates the e200 OnCE register addresses. Attempted writes to read-only registers are ignored.

Table 8-8 indicates the e200 OnCE register addresses.

Table 8-8. e200 OnCE Register Addressing

RS[0:6]	Register Selected
000 0000	Reserved
000 0001	Reserved
000 0010	JTAG ID (read-only)
000 0011–000 1111	Reserved
001 0000	CPU Scan Register (CPUSCR)
001 0001	No Register Selected (Bypass)
001 0010	OnCE Control Register (OCR)
001 0011	Reserved
001 0100–001 1111	Reserved
010 0000	Instruction Address Compare 1 (IAC1)
010 0001	Instruction Address Compare 2 (IAC2)
010 0010	Instruction Address Compare 3 (IAC3)
010 0011	Instruction Address Compare 4 (IAC4)
010 0100	Data Address Compare 1 (DAC1)
010 0101	Data Address Compare 2 (DAC2)
010 0110	Reserved (DVC1 future use)
010 0111	Reserved (DVC2 future use)
010 1000–010 1011	Reserved
010 1100	Reserved (DBCNT)
010 1101–010 1111	Reserved
011 0000	Debug Status Register (DBSR)
011 0001	Debug Control Register 0 (DBCR0)
011 0010	Debug Control Register 1 (DBCR1)
011 0011	Debug Control Register 2 (DBCR2)
011 0100–101 1111	Reserved (do not access)
110 0000–110 1110	Reserved (do not access)
110 1111	Shared Nexus Control Register Register Select
111 0000–111 1001	General Purpose register selects [0:9]
111 1010	(Reserved))
111 1011	(Reserved)
111 1100	Nexus2/3-Access
111 1101	Reserved

Table 8-8. e200 OnCE Register Addressing (continued)

RS[0:6]	Register Selected
111 1110	Enable_OnCE ¹
111 1111	Bypass

¹ Causes assertion of the j_en_once_regssel output. Refer to [Section 6.3.15.15](#), “Enable Once Register Select (j_en_once_regssel).”

The Once Decoder receives as input the 10-bit command from the OCMD, and status signals from the processor, and generates all the strobes required for reading and writing the selected OnCE registers.

Single stepping of instructions is performed by placing the CPU in debug mode, scanning in appropriate information into the CPUSCR, and setting the Go bit (with the EX bit cleared) with the RS field indicating either the CPUSCR or No Register Selected. After executing a single instruction, the CPU re-enters debug mode and await further commands. During single-stepping, exception conditions may occur if not properly masked by debug firmware (interrupts, machine checks, bus error conditions, etc.) and may prevent the desired instruction from being successfully executed. The OSR[ERR] bit is set to indicate this condition. In these cases, values in the CPUSCR correspond to the first instruction of the exception handler.

Additionally, the DBCR0[EDM] bit is forced to ‘1’ internally while single-stepping to prevent Debug events from generating Debug interrupts. Also, during a debug session, the DBSR is frozen from updates due to debug events regardless of DBCRO[EDM]. They may still be modified during a debug session via a single-stepped **mtspr** instruction if DBCRO[EDM] is programmed to a ‘0’, or via OnCE access if DBCR0[EDM] is set.

8.4.5.3 e200 OnCE Control Register (OCR)

The e200 OnCE Control Register is a 32-bit register used to force the e200 core into debug mode and to enable / disable sections of the e200 OnCE control logic. It also provides control over the MMU during a debug session. The control bits are read/write. These bits are only effective while OnCE is enabled (*jd_en_once* asserted). The OCR is shown in [Figure 8-10](#).




Figure 8-10. OnCE Control Register

Table 8-9 provides bit definitions for the OnCE Control Register.

Table 8-9. OnCE Control Register Bit Definitions

Bit(s)	Name	Description
0:7	—	Reserved
8	I_DMDIS ¹	Instruction Side Debug MMU Disable Control Bit (I_DMDIS) 0 MMU not disabled for debug sessions 1 MMU disabled for debug sessions This bit may be used to control whether the MMU is enabled normally, or whether the MMU is disabled during a debug session for Instruction Accesses. When enabled, the MMU functions normally. When disabled, for Instruction Accesses, no address translation is performed (1:1 address mapping), and the TLB VLE, I,M, and E bits are taken from the OCR bits I_VLE, I_DI, I_DM, and I_DE bits. The W and G bits are assumed '0'. The SX and UX access permission control bits are set to '1' to allow full access. When disabled, no TLB miss or TLB exceptions are generated for Instruction accesses. External access errors can still occur.
9:10	—	Reserved
11	I_DVLE ¹	Instruction Side Debug TLB 'VLE' Attribute Bit (I_DVLE) This bit is used to provide the 'VLE' attribute bit to be used when the MMU is disabled during a debug session.
12	I_DI ¹	Instruction Side Debug TLB 'I' Attribute Bit (I_DI) This bit is used to provide the 'I' attribute bit to be used for Instruction accesses when the MMU is disabled for Instruction accesses during a debug session.
13	I_DM ¹	Instruction Side Debug TLB 'M' Attribute Bit (I_DM) This bit is used to provide the 'M' attribute bit to be used for Instruction accesses when the MMU is disabled for Instruction accesses during a debug session.
14	—	Reserved
15	I_DE ¹	Instruction Side Debug TLB 'E' Attribute Bit (I_DE) This bit is used to provide the 'E' attribute bit to be used for Instruction accesses when the MMU is disabled for Instruction accesses during a debug session.
16	D_DMDIS ¹	Data Side Debug MMU Disable Control Bit (D_DMDIS) 0 MMU not disabled for debug sessions 1 MMU disabled for debug sessions This bit may be used to control whether the MMU is enabled normally, or whether the MMU is disabled during a debug session for Data Accesses. When enabled, the MMU functions normally. When disabled, for Data Accesses, no address translation is performed (1:1 address mapping), and the TLB WIMGE bits are taken from the OCR bits D_DW, D_DI, D_DM, D_DG, and D_DE bits. The SR, SW, UR, and UW access permission control bits are set to '1' to allow full access. When disabled, no TLB miss or TLB exceptions are generated for Data accesses. External access errors can still occur.
17:18	—	Reserved
19	D_DW ¹	Data Side Debug TLB 'W' Attribute Bit (D_DW) This bit is used to provide the 'W' attribute bit to be used for Data accesses when the MMU is disabled for Data accesses during a debug session.
20	D_DI ¹	Data Side Debug TLB 'I' Attribute Bit (D_DI) This bit is used to provide the 'I' attribute bit to be used for Data accesses when the MMU is disabled for Data accesses during a debug session.
21	D_DM ¹	Data Side Debug TLB 'M' Attribute Bit (D_DM) This bit is used to provide the 'M' attribute bit to be used for Data accesses when the MMU is disabled for Data accesses during a debug session.

Table 8-9. OnCE Control Register Bit Definitions (continued)

Bit(s)	Name	Description
22	D_DG ¹	Data Side Debug TLB 'G' Attribute Bit (D_DG) This bit is used to provide the 'G' attribute bit to be used for Data accesses when the MMU is disabled for Data accesses during a debug session.
23	D_DE ¹	Data Side Debug TLB 'E' Attribute Bit (D_DE) This bit is used to provide the 'E' attribute bit to be used for Data accesses when the MMU is disabled for Data accesses during a debug session.
24:28	—	Reserved
29	WKUP	Wakeup Request Bit (WKUP) This control bit may be used to force the e200 <i>p_wakeup</i> output signal to be asserted. This control function may be used by debug firmware to request that the chip-level clock controller restore the <i>m_clk</i> input to normal operation regardless of whether the CPU is in a low power state to ensure that debug resources may be properly accessed by external hardware through scan sequences.
30	FDB	Force Breakpoint Debug Mode Bit (FDB) This control bit is used to determine whether the processor is operating in breakpoint debug enable mode or not. The processor may be placed in breakpoint debug enable mode by setting this bit. In breakpoint debug enable mode, execution of the ' <i>bkpt</i> ' pseudo- instruction causes the processor to enter debug mode, as if the ' <i>jd_de_b</i> ' input had been asserted. This bit is qualified with DBCR0[EDM], which must be set for FDB to take effect.
31	DR	CPU Debug Request Control Bit This control bit is used to unconditionally request the CPU to enter the Debug Mode. The CPU indicates that Debug Mode has been entered via the data scanned out in the shift-IR state. 0 No Debug Mode request 1 Unconditional Debug Mode request When the DR bit is set the processor enters Debug mode at the next instruction boundary.

¹ Unused by Zen Z0n2p and Zen Z0Hn2p

8.4.6 Access to Debug Resources

Resources contained in the e200 OnCE Module which do not require the e200 processor core to be halted for access may be accessed while the e200 core is running, and does not interfere with processor execution. Accesses to other resources such as the CPUSCR require the e200 core to be placed in debug mode to avoid synchronization hazards. Debug firmware may ensure that it is safe to access these resources by determining the state of the e200 core prior to access. Note that a scan operation to update the CPUSCR is required prior to exiting debug mode if debug mode has been entered.

Some cases of write accesses other than accesses to the OnCE Command and Control registers, or the EDM bit of DBCR0 require the e200 *m_clk* to be running for proper operation. The OnCE control register provides a means of signaling this need to a system level clock control module.

In addition, because the CPU may cause multiple bits of certain registers to change state, reads of certain registers while the CPU is running (DBSR, etc.) may not have consistent bit settings unless read twice with the same value indicated. In order to guarantee that the contents are consistent, the CPU should be placed into debug mode, or multiple reads should be performed until consistent values have been obtained on consecutive reads.

[Table 8-10](#) provides a list of access requirements for OnCE registers.

Table 8-10. OnCE Register Access Requirements

Register Name	Access Requirements					Notes
	Requires <i>jd_en_once</i> to be Asserted	Requires DBCR0 [EDM] = 1	Requires <i>m_clk</i> Active for Write Access	Requires CPU to be Halted for Read Access	Requires CPU to be Halted for Write Access	
Enable_OnCE	N	N	N	N	—	—
Bypass	N	N	N	N	N	—
CPUSCR	Y	Y	Y	Y	Y	—
DAC1	Y	Y	Y	N	*1	—
DAC2	Y	Y	Y	N	*1	—
DBCR0	Y	Y	Y	N	*1	*DBCR0[EDM] access only requires <i>jd_en_once</i> asserted
DBCR1	Y	Y	Y	N	*1	—
DBCR2	Y	Y	Y	N	*1	—
DBSR	Y	Y	Y	N ²	*1	—
IAC1	Y	Y	Y	N	*1	—
IAC2	Y	Y	Y	N	*1	—
IAC3	Y	Y	Y	N	*1	—
IAC4	Y	Y	Y	N	*1	—
JTAG ID	N	N	—	N	—	Read-only
OCR	Y	N	N	N	N	—
OSR	Y	N	—	N	—	Read-only, accessed by scanning out IR while <i>jd_en_once</i> is asserted
Cache Debug Access Control (CDACNTL) ²	Y	N	Y	Y	Y	CPU must be in debug mode with clocks running
Cache Debug Access Data (CDADATA) ²	Y	N	Y	Y	Y	CPU must be in debug mode with clocks running
Nexus2/3-Access	Y	N	N	N	N	—
External GPRs	Y	N	N	N	N	—
LSRL Select	Y	N	?	?	?	System Test logic implementation determines LSRL functionality

¹ Writes to these registers while the CPU is running may have unpredictable results due to the pipelined nature of operation, and the fact that updates are not synchronized to a particular clock, instruction, or bus cycle boundary, therefore it is strongly recommended to ensure the processor is first placed into debug mode before updates to these registers are performed.

² Not present on Zen Z0n2p or Zen Z0Hn2p

8.4.7 Methods of Entering Debug Mode

The OnCE Status Register indicates that the CPU has entered the debug mode via the DEBUG status bit. The following sections describe how e200 Debug Mode is entered assuming the OnCE circuitry has been enabled. e200 OnCE operation is enabled by the assertion of the *jd_en_once* input (see Section 8.4.4.1).

8.4.7.1 External Debug Request During RESET

Holding the *jd_de_b* signal asserted during the assertion of *p_reset_b*, and continuing to hold it asserted following the negation of *p_reset_b* causes the e200 core to enter Debug Mode. After receiving an acknowledge via the OnCE Status Register DEBUG bit, the external command controller should negate the *jd_de_b* signal before sending the first command. Note that in this case the e200 core does not execute an instruction before entering Debug Mode, although the first instruction to be executed may be fetched prior to entering Debug Mode.

In this case, all values in the debug scan chain are undefined, and the external Debug Control Module is responsible for proper initialization of the chain before debug mode is exited. In particular, the exception processing associated with reset, may not be performed when the debug mode is exited, thus, the Debug controller must initialize the PC, MSR, and IR to the image that the processor would have obtained in performing reset exception processing, or must cause the appropriate reset to be re-asserted.

8.4.7.2 Debug Request During RESET

Asserting a debug request by setting the DR bit in the OCR during the assertion of *p_reset_b* causes the chip to enter debug mode. In this case the chip may fetch the first instruction of the reset exception handler, but does not execute an instruction before entering debug mode. In this case, all values in the debug scan chain are undefined, and the external Debug Control Module is responsible for proper initialization of the chain before debug mode is exited. In particular, the exception processing associated with reset may not be performed when the debug mode is exited, thus, the Debug controller must initialize the PC, MSR, and IR to the image that the processor would have obtained in performing reset exception processing, or must cause the appropriate reset to be re-asserted.

8.4.7.3 Debug Request During Normal Activity

Asserting a debug request by setting the DR bit in the OCR during normal chip activity causes the chip to finish the execution of the current instruction and then enter the debug mode. Note that in this case the chip completes the execution of the current instruction and stops after the newly fetched instruction enters the CPU instruction register. This process is the same for any newly fetched instruction including instructions fetched by the interrupt processing, or those that are aborted by the interrupt processing.

8.4.7.4 Debug Request During Waiting, Halted or Stopped State

Asserting a debug request by setting the DR bit in the OCR when the chip is in the Waiting state (*p_waiting* asserted), Halted state (*p_halted* asserted) or Stopped state (*p_stopped* asserted) causes the CPU to exit the state and enter the debug mode once the CPU clock *m_clk* has been restored. Note that in this case, the

CPU negates the *p_waiting*, *p_halted* and *p_stopped* outputs. Once the debug session has ended, the CPU returns to the state it was in prior to entering debug mode.

To signal the chip-level clock generator to re-enable *m_clk*, the *p_wakeup* output is asserted whenever the debug block is asserting a debug request to the CPU due to OCR[DR] being set, or *jd_de_b* assertion, and remains set from then until the debug session ends (*jd_debug_b* goes from asserted to negated). In addition, the status of the *jd_mclk_on* input (after synchronization to the *j_tclk* clock domain) may be sampled along with other status bits from the *j_tdo* output during the Shift_IR TAP controller state. This status may be used if necessary by external debug firmware to ensure proper scan sequences occur to registers in the *m_clk* clock domain.

8.4.7.5 Software Request During Normal Activity

Upon executing a “*bkpt*” pseudo-instruction (for e200, defined to be an all 0’s instruction opcode) when the OCR register’s (FDB) bit is set (debug mode enable control bit is true), and DBCR0[EDM]=1, the CPU enters the debug mode after the instruction following the “*bkpt*” pseudo-instruction has entered the instruction register.

8.4.8 CPU Status and Control Scan Chain Register (CPUSCR)

A number of on-chip registers store the CPU pipeline status and are configured in a single scan chain for access by the e200 OnCE controller. The CPUSCR register contains these processor resources, which are used to restore the pipeline and resume normal chip activity upon return from the debug mode, as well as a mechanism for the emulator software to access processor and memory contents. [Figure 8-11](#) shows the block diagram of the pipeline information registers contained in the CPUSCR. Once debug mode has been entered, it is required to scan in and update this register prior to exiting debug mode.




Figure 8-11. CPU Scan Chain Register (CPUSCR)

8.4.8.1 Instruction Register (IR)

The Instruction Register (IR) provides a mechanism for controlling the debug session by serving as a means for forcing in selected instructions, and then causing them to be executed in a controlled manner by the debug control block. The opcode of the next instruction to be executed when entering debug mode is contained in this register when the scan-out of this chain begins. This value should be saved for later restoration if continuation of the normal instruction stream is desired.

On scan-in, in preparation for exiting debug mode, this register is filled with an instruction opcode selected by debug control software. By selecting appropriate instructions and controlling the execution of those instructions, the results of execution may be used to examine or change memory locations and processor registers. The debug control module external to the processor core controls execution by providing a single-step capability. Once the debug session is complete and normal processing is to be resumed, this register may be loaded with the value originally scanned out.

8.4.8.2 Control State Register (CTL)

The Control State Register (CTL) is a 32-bit register that stores the value of certain internal CPU state variables before the debug mode is entered. This register is affected by the operations performed during the debug session and should normally be restored by the external command controller when returning to normal mode. In addition to saved internal state variables, two of the bits are used by emulation firmware to control the debug process. In certain circumstances, emulation firmware must modify the content of this register as well as the PC and IR values in the CPUSCR before exiting debug mode. These cases are described below. [Figure 8-12](#).




Figure 8-12. Control State Register (CTL)

WAITING—WAITING State Status

This bit indicates whether the CPU was in the waiting state prior to entering debug mode. If set, the CPU was in the waiting state. Upon exiting a debug session, the value of this bit in the restored CPUSCR determines whether the CPU re-enters the waiting state on a go+exit.

- 0—CPU was not in the waiting state when debug mode was entered
- 1—CPU was in the waiting state when debug mode was entered

PCOFST—PC Offset Field

This field indicates whether the value in the PC portion of the CPUSCR must be adjusted prior to exiting debug mode. Due to the pipelined nature of the CPU, the PC value must be backed-up by emulation software in certain circumstances. The PCOFST field specifies the value to be subtracted from the original value of the PC. This adjusted PC value should be restored into the PC portion of the CPUSCR just prior to exiting debug mode with a go+exit. In the event the PCOFST is non-zero, the IR should be loaded with a nop instruction instead of the original IR value, other wise the original value of IR should be restored. (But see PCINV which overrides this field)

- 0000—No correction required.
- 0001—Subtract 0x04 from PC.
- 0010—Subtract 0x08 from PC.
- 0011—Subtract 0x0C from PC.
- 0100—Subtract 0x10 from PC.
- 0101—Subtract 0x14 from PC.
- All other encodings are reserved

* — Internal State Bits

These control bits represent internal processor state and should be restored to their original value after a debug session is completed, i.e when a e200 OnCE command is issued with the GO and EX bits set and not ignored. When performing instruction execution during a debug

session (see [Section 8.4.4.3, “e200 OnCE Debug Output \(jd_debug_b\)”](#)) which is not part of the normal program execution flow, these bits should be set to a 0.

PCINV—PC and IR Invalid Status Bit

This status bit indicates that the values in the IR and PC portions of the CPUSCR are invalid. Exiting debug mode with the saved values in the PC and IR have unpredictable results. Debug firmware should initialize the PC and IR values in the CPUSCR with desired values prior to exiting debug mode if this bit was set when debug mode was initially entered.

- 0 = No error condition exists.
- 1 = Error condition exists. PC and IR are corrupted.

FFRA—Feed Forward RA Operand Bit

This control bit causes the content of the WBBR_{low} to be used as the RA (RS for logical and shift operations or RX for VLE se_ instructions) operand value of the first instruction to be executed following an update of the CPUSCR. This allows the debug firmware to update processor registers — initialize the WBBR_{low} with the desired value, set the FFRA bit, and execute a ori Rx,Rx,0 instruction to the desired register.

- 0 = No action.
- 1 = Content of WBBR_{low} used as RA (RS for logical and shift operations) operand value

IRStat0—IR Status Bit 0

This control bit indicates a TEA status for the IR.

- 0 = No TEA occurred on the fetch of this instruction.
- 1 = TEA occurred on the fetch of this instruction.

IRStat1—IR Status Bit 1

This control bit indicates a TLB Miss status for the IR. (Note: this bit is reserved.)

- 0 = No TLB Miss occurred on the fetch of this instruction.
- 1 = TLB Miss occurred on the fetch of this instruction.

IRStat2—IR Status Bit 2

This control bit indicates an Instruction Address Compare 1 event status for the IR.

- 0 = No Instruction Address Compare 1 event occurred on the fetch of this instruction.
- 1 = An Instruction Address Compare 1 event occurred on the fetch of this instruction.

IRStat3—IR Status Bit 3

This control bit indicates an Instruction Address Compare 2 event status for the IR.

- 0 = No Instruction Address Compare 2 event occurred on the fetch of this instruction.
- 1 = An Instruction Address Compare 2 event occurred on the fetch of this instruction.

IRStat4—IR Status Bit 4

This control bit indicates an Instruction Address Compare 3 event status for the IR.

- 0 = No Instruction Address Compare 3 event occurred on the fetch of this instruction.
- 1 = An Instruction Address Compare 3 event occurred on the fetch of this instruction.

IRStat5—IR Status Bit 5

This control bit indicates an Instruction Address Compare 4 event status for the IR.

- 0 = No Instruction Address Compare 4 event occurred on the fetch of this instruction.
- 1 = An Instruction Address Compare 4 event occurred on the fetch of this instruction.

IRStat6—IR Status Bit 6

This control bit indicates a Parity Error status for the IR. (Note: this bit is reserved.)

- 0 = No Parity Error occurred on the fetch of this instruction.
- 1 = Parity Error occurred on the fetch of this instruction.

IRStat7—IR Status Bit 7

This control bit indicates a Precise External Termination Error status for the IR.

- 0 = 0 = No Precise External Termination Error occurred on the fetch of this instruction.
- 1 = Precise External Termination Error occurred on the fetch of this instruction.

IRStat8—IR Status Bit 8

This control bit indicates the Power Architecture VLE status for the IR. (Note: this bit is always set on Zen Z0n2p and Zen Z0Hn2p.)

- 0 = IR contains a Book E instruction.
- 1 = IR contains a Power Architecture VLE instruction, aligned in the Most Significant Portion of IR if 16-bit.

IRStat9—IR Status Bit 9

This control bit indicates the Power Architecture VLE Byte-ordering Error status for the IR, or a Book E misaligned instruction fetch, depending on the state of IRStat8. (Note: this bit is reserved on Zen Z0n2p and Zen Z0Hn2p.)

- 0 = IR contains an instruction without a byte-ordering error and no Misaligned Instruction Fetch Exception has occurred (no MIF).
- 1 = If IRStat8 = '0', A Book E Misaligned Instruction Fetch Exception has occurred while filling the IR.
If IRStat8 = '1', IR contains an instruction with a byte-ordering error due to mismatched VLE page attributes, or due to E indicating little-endian for a VLE page.

Emulation firmware should modify the content of the CTL, PC, and IR values in the CPUSCR during execution of debug related instructions as well as just prior to exiting debug with a go+exit command. During the debug session, the CTL register should be written with the FFRA bit set as appropriate, and all other bit set to '0', and the IR set to the value of the desired instruction to be executed. IRStat8 is used to determine the type of instruction present in the IR.

Just prior to exiting debug mode with a go+exit, the PCINV status bit which was originally present when debug mode was first entered should be tested, and if set, the PC and IR initialized for performing whatever recovery sequence is appropriate for a faulted exception vector fetch. If the PCINV bit is cleared, then the

PCOFST bits should be examined to determine whether the PC value must be adjusted. Due to the pipelined nature of the CPU, the PC value must be backed-up by emulation software in certain circumstances. The PCOFST field specifies the value to be subtracted from the original value of the PC. This adjusted PC value should be restored in to the PC portion of the CPUSCR just prior to exiting debug mode with a `go+exit`. In the event the PCOFST is non-zero, the IR should be loaded with a `nop` instruction (such as `ori r0,r0,0`) instead of the original IR value, otherwise the original value of IR should be restored. Note that when a correction is made to the PC value, it generally points to the last completed instruction, although that instruction is not re-executed. The `nop` instruction is executed instead, and instruction fetch and execution resumes at location PC+4. IRStat8 is used to determine the type of instruction present in the IR, thus should be cleared in this case.

For the CTL register, the internal state bits should be restored to their original value. The IRStatus bits should be set to '0's if the PC was adjusted. If no PC adjustment was performed, emulation firmware should determine whether IRStat2-5 should be set to '0' to avoid re-entry into debug mode for an instruction breakpoint request. Upon exiting debug mode with `go+exit`, if one of these bits is set, debug mode is re-entered prior to any further instruction execution.

8.4.8.3 Program Counter Register (PC)

The PC is a 32-bit register that stores the value of the program counter which was present when the chip entered the debug mode. It is affected by the operations performed during the debug mode and must be restored by the external command controller when the CPU returns to normal mode. PC normally points to the instruction contained in the IR portion of CPUSCR. If debug firmware wishes to redirect program flow to an arbitrary location, the PC and IR should be initialized to correspond to the first instruction to be executed upon resumption of normal processing. Alternatively, the IR may be set to a `nop` and the PC set to point to the location prior to the location at which it is desired to redirect flow to. On exiting debug mode, the `nop` is executed, and instruction fetch and execution resumes at PC+4.

8.4.8.4 Write-Back Bus Register (WBBR_{low}, WBBR_{high})

WBBR is used as a means of passing operand information between the CPU and the external command controller. Whenever the external command controller needs to read the contents of a register or memory location, it forces the chip to execute an instruction that brings that information to WBBR. WBBR_{low} holds the 32-bit result of most instructions including load data returned for a load or load with update instruction. WBBR_{high} holds the updated effective address calculated by a load with update instruction. It is undefined for other instructions.

As an example, to read the lower 32 bits of processor register `r1`, an `e_ori r1,r1,0` instruction is executed, and the result value of the instruction is latched into WBBR_{low}. The contents of WBBR_{low} can then be delivered serially to the external command controller. To update a processor resource, this register is initialized with a data value to be written, and an `e_ori` instruction is executed which uses this value as a substitute data value. The Control State register FFRA bit forces the value of the WBBR_{low} to be substituted for the normal RS source value of the `e_ori` instruction, thus allowing updates to processor registers to be performed (refer to [Section 8.4.8.2, "Control State Register \(CTL\),"](#) for more detail on the CTL[FFRA] bit).

WBBR_{low} and WBBR_{high} are generally undefined on instructions which do not writeback a result, and due to control issues are not defined on **lmw** or branch instructions as well.

8.4.8.5 Machine State Register (MSR)

The MSR is a 32-bit register used to read/write the Machine State Register. Whenever the external command controller needs to save or modify the contents of the Machine State Register, this register is used. This register is affected by the operations performed during the debug mode and must be restored by the external command controller when returning to normal mode.

8.4.9 Reserved Registers (Reserved)

The Reserved Registers are used to control various test control logic. These registers are not intended for customer use. To preclude device and/or system damage, these registers should not be accessed.

8.5 Watchpoint Support

e200 supports the generation and signalling of watchpoints when operating in internal debug mode (DBCR0[IDM]=1) or in external debug mode (DBCR0[EDM]=1). Watchpoints are indicated with a dedicated set of interface signals. The *jd_watchpoint*[0:5] output signals are used to indicate that a watchpoint has occurred.

Each debug address compare function (IAC1-4, DAC1-2) is capable of triggering a watchpoint output. The DBCRx control fields are used to configure watchpoints, regardless of whether events are enabled in DBCR0. Watchpoints may occur whenever an associated event would have been posted in the Debug Status Register if enabled. No explicit enable bits are provided for watchpoints; they are always enabled by definition (except during a debug session). If not desired, the base address values for these events may be programmed to an unused system address. MSR[DE] has no effect on watchpoint generation.

External logic may monitor the assertion of these signals for debugging purposes. Watchpoints are signaled in the clock cycle following the occurrence of the actual event. The Nexus2+ module also monitors assertion of these signals for various development control purposes.

Table 8-11. Watchpoint Output Signal Assignments

Signal Name	Type	Description
<i>jd_watchpt</i> [0]	IAC1	Instruction Address Compare 1 watchpoint Asserted whenever an IAC1 compare occurs regardless of being enabled to set DBSR status
<i>jd_watchpt</i> [1]	IAC2	Instruction Address Compare 2 watchpoint Asserted whenever an IAC2 compare occurs regardless of being enabled to set DBSR status
<i>jd_watchpt</i> [2]	IAC3	Instruction Address Compare 3 watchpoint Asserted whenever an IAC3 compare occurs regardless of being enabled to set DBSR status
<i>jd_watchpt</i> [3]	IAC4	Instruction Address Compare 4 watchpoint Asserted whenever an IAC4 compare occurs regardless of being enabled to set DBSR status

Table 8-11. Watchpoint Output Signal Assignments (continued)

Signal Name	Type	Description
<i>jd_watchpt[4]</i>	DAC1 ¹	Data Address Compare 1 watchpoint Asserted whenever a DAC1 compare occurs regardless of being enabled to set DBSR status
<i>jd_watchpt[5]</i>	DAC2 ¹	Data Address Compare 2 watchpoint Asserted whenever a DAC2 compare occurs regardless of being enabled to set DBSR status

¹ If the corresponding event is completely disabled in DBCR0, either load-type or store-type data accesses are allowed to generate watchpoints, otherwise watchpoints are generated only for the enabled conditions.

8.6 Basic Steps for Enabling, Using, and Exiting External Debug Mode

The following steps show one possible scenario for a debugger wishing to use the external debug facilities. This simplified flow is intended to illustrate basic operations, but does not cover all potential methods in depth.

Enabling External Debug Mode and initializing Debug registers

- The debugger should ensure that the *jd_en_once* control signal is asserted in order to enable OnCE operation
- Select the OCR and write a value to it in which OCR[DR], OCR[WKUP], are set to '1'. The tap controller must step through the proper states as outlined earlier. This step places the CPU in a debug state in which it is halted and awaiting single-step commands or a release to normal mode
- Scan out the value of the OSR to determine that the CPU clock is running and the CPU has entered the Debug state. This can be done in conjunction with a Read of the CPUSCR. The OSR is shifted out during the Shift_IR state. The CPUSCR is shifted out during the Shift_DR state. The debugger should save the scanned-out value of CPUSCR for later restoration.
- Select the DBCR0 register and update it with the DBCR0[EDM] bit set
- Clear the DBSR status bits
- Write appropriate values to the DBCRx, IAC, DAC registers. Note that the initial write to DBCR0 only affects the EDM bit, so the remaining portion of the register must now be initialized, keeping the EDM bit set

At this point the system is ready to commence debug operations. Depending on the desired operation, different steps must occur.

- Optionally, set the OCR[I_DMDIS] and/or OCR[I_DMDIS] control bits to ensure that no TLB misses occur while performing the debug operations
- Optionally, ensure that the values entered into the MSR portion of the CPUSCR during the following steps cause interrupt to be disabled (clearing MSR[EE] and MSR[CE]). This ensures that external interrupt sources do not cause single-step errors.

To single-step the CPU:

- debugger scans in either a new or a previously saved value of the CPUSCR (with appropriate modification of the PC and IR as described in [Section 8.4.8.2, “Control State Register \(CTL\)”](#)), with a Go+Noexit OnCE Command value.
- The debugger scans out the OSR with ‘no-register selected’, Go cleared, and determines that the PCU has re-entered the Debug state and that no ERR condition occurred

To return the CPU to normal operation (without disabling external debug mode)

- The OCR[DMDIS], OCR[DR], control bits should be cleared, leaving the OCR[WKUP] bit set
- The debugger restores the CPUSCR with a previously saved value of the CPUSCR (with appropriate modification of the PC and IR as described in [Section 8.4.8.2, “Control State Register \(CTL\)”](#)), with a Go+Exit OnCE Command value.
- The OCR[WKUP] bit may then be cleared

To exit External Debug Mode

- The debugger should place the CPU in the debug state via the OCR[DR] with OCR[WKUP] asserted, scanning out and saving the CPUSCR
- The debugger should write the DBCRx registers as needed, likely clearing every enable except the DBCR0[EDM] bit
- The debugger should write the DBSR to a cleared state
- The debugger should re-write the DBCR0 with all bits including EDM cleared
- The debugger should clear the OCR[DR] bit
- The debugger restores the CPUSCR with the previously saved value of the CPUSCR (with appropriate modification of the PC and IR as described in [Section 8.4.8.2, “Control State Register \(CTL\)”](#)), with a Go+Exit OnCE Command value.
- The OCR[WKUP] bit may then be cleared

NOTE

These steps are meant by way of examples, and are not meant to be an exact template for debugger operation.

Chapter 9

Nexus 2+ Module

The e200z0 and e200z0h Nexus 2+ module provides real-time development capabilities for Zen processors in compliance with the IEEE-ISTO Nexus 5001-2003. This module provides development support capabilities without requiring the use of address and data pins for internal visibility.

A portion of the pin interface (the JTAG port) is also shared with the OnCE/Nexus 1 unit. IEEE-ISTO 5001-2003 defines an extensible auxiliary port which is used in conjunction with the JTAG port in Zen processors.

The Nexus modules are coupled to the core and monitor a variety of signals including addresses, data, control signals, status signals, etc. In some SoC designs, there may be a single shared Nexus module with the capability of selectively monitoring more than one CPU. Control over this selection of the source if information is provided by a SoC-level Shared Nexus Control Module, which is accessed through JTAG via the Nexus1 Shared Nexus Control register. Specifics of this module are provided in a separate document. The CPU provides an interface signal to communicate selection of this register.

9.1 Introduction

9.1.1 General Description

This chapter defines the auxiliary pin functions, transfer protocols and standard development features of a Class 2 device in compliance with the IEEE-ISTO Nexus 5001-2003. The development features supported are Program Trace, Data Trace, Watchpoint Messaging, Ownership Trace, and Read/Write Access via the JTAG interface. The Nexus 2+ module also supports two Class 4 features: Watchpoint Triggering and Processor Overrun Control.

9.1.2 Terms and Definitions

Table 9-1 contains a set of terms and definitions associated with the Nexus 2+ module.

Table 9-1. Terms and Definitions

Term	Description
IEEE-ISTO 5001	Consortium and standard for real-time embedded system design. World wide Web documentation at http://www.ieee-isto.org/Nexus5001
Auxiliary Port	Refers to Nexus auxiliary port. Used as auxiliary port to the IEEE 1149.1 JTAG interface.
Branch Trace Messaging (BTM)	Visibility of addresses for taken branches and exceptions, and the number of sequential instructions executed between each taken branch.

Table 9-1. Terms and Definitions (continued)

Term	Description
JTAG Compliant	Device complying to IEEE 1149.1 JTAG standard
JTAG IR and DR Sequence	JTAG Instruction Register (IR) scan to load an opcode value for selecting a development register. The JTAG IR corresponds to the OnCE command register (OCMD). The selected development register is then accessed via a JTAG Data Register (DR) scan.
Nexus1	The Zen (OnCE) debug module. This module integrated with each Zen processor provides all static (core halted) debug functionality. This module is compliant with Class1 of IEEE-ISTO 5001.
Ownership Trace Message (OTM)	Visibility of process/function that is currently executing.
Public Messages	Messages on the auxiliary pins for accomplishing common visibility and controllability requirements
SOC	“System-on-a-Chip”. SOC signifies all of the modules on a single die. This generally includes one or more processors with associated peripherals, interfaces and memory modules.
Standard	The phrase “according to the standard” is used to indicate according to IEEE-ISTO 5001.
Transfer Code (TCODE)	Message header that identifies the number and/or size of packets to be transferred, and how to interpret each of the packets.
Watchpoint	A Data or Instruction Breakpoint which does not cause the processor to halt. Instead, a pin is used to signal that the condition occurred. A Watchpoint Message is also generated.

9.1.3 Feature List

The Nexus 2+ module is compatible with Class 2 of IEEE-ISTO 5001-2003, with additional Class 3 and Class 4 features available. The following features are implemented:

- Program Trace via Branch Trace Messaging (BTM). Branch trace messaging displays program flow discontinuities (direct and indirect branches, exceptions, etc.), allowing the development tool to interpolate what transpires between the discontinuities. Thus static code may be traced.
- Ownership Trace via Ownership Trace Messaging (OTM). OTM facilitates ownership trace by providing visibility of which process ID or operating system task is activated. An Ownership Trace Message is transmitted when a new process/task is activated, allowing the development tool to trace ownership flow.
- Run-time access to embedded processor memory map via the JTAG port. This allows for enhanced download/upload capabilities.
- Watchpoint Messaging via the auxiliary pins
- Watchpoint Trigger enable of Program and/or Data Trace Messaging
- Auxiliary interface for higher data input/output
 - Configurable (min/max) Message Data Out pins (*nex_mdo[n:0]*)
 - One (1) or two (2) Message Start/End Out pins (*nex_mseo_b[1:0]*)

- One (1) Read/Write Ready pin (*nex_rdy_b*) pin
- One (1) Watchpoint Event pin (*nex_evto_b*)
- One (1) Event In pin (*nex_evti_b*)
- One (1) MCKO (Message Clock Out) pin
- Registers for Program Trace, Ownership Trace and Watchpoint Trigger.
- All features controllable and configurable via the JTAG port

NOTE

For multi-Nexus implementations, the configuration of the Message Data Out pins is controlled by the Port Control Register (@ the SoC level). For single Nexus implementations, this configuration is controlled by Development Control Register 1 (DC1) within the e200 Nexus 2+ module.

In either implementation, Full Port Mode (FPM—maximum number of MDO pins) or Reduced Port Mode (RPM—minimum number of MDO pins) are supported. This setting should not be changed while the system is running.

NOTE

The configuration of the Message Start/End Out pins (1 or 2) is determined at the SOC integration level. This option is hard-wired based on SOC bandwidth requirements.

9.1.4 Functional Block Diagram




Figure 9-1. Nexus 2+ Functional Block Diagram

9.2 Enabling Nexus 2+ Operation

The Nexus module is enabled by loading a single instruction (*NEXUS2-ACCESS*) into the JTAG Instruction Register (IR) (OnCE OCMD register). For the e200 Nexus 2+ module, the OCMD value is 0b0001111100. Once enabled, the module is ready to accept control input via the JTAG/OnCE pins.

Enabling the Nexus 2+ module automatically enables the generation of Debug Status Messages.

The Nexus module is disabled when the JTAG state machine reaches the Test-Logic-Reset state. This state can be reached by the assertion of the *j_trst_b* pin or by cycling through the state machine using the *j_tms* pin. The Nexus module also is disabled if a Power-on-Reset (POR) event occurs. If the Nexus 2+ module is disabled, no trace output is provided, and the module disables (drives inactive) auxiliary port output pins (*nex_mdo*[n:0], *nex_mseo*[1:0], *nex_mcko*). Nexus registers are not available for reads or writes.

NOTE

Please refer to the “Nexus 2+ Integration Guide” for details on IEEE-ISTO 5001 compatibility w.r.t. output pins and multiple Nexus module configurations.

9.3 TCODEs Supported

The Nexus 2+ pins allow for flexible transfer operations via Public Messages. A TCODE defines the transfer format, the number and/or size of the packets to be transferred, and the purpose of each packet. IEEE-ISTO 5001-2003 defines a set of public messages. The Nexus 2+ block supports the public TCODEs seen in [Table 9-2](#). Each message contains multiple packets transmitted in the order shown in the table.

Table 9-2. Public TCODEs Supported

Message Name	Minimum Packet Size (bits)	Maximum Packet Size (bits)	Packet Type	Packet Description
Debug Status	6	6	fixed	TCODE number = 0
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	8	8	fixed	Debug Status Register (DS[31:24])
Ownership Trace Message	6	6	fixed	TCODE number = 2
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	32	32	fixed	Task/Process ID tag
Program Trace—Direct Branch Message	6	6	fixed	TCODE number = 3
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	1	8	variable	# sequential instructions executed since last taken branch
Program Trace—Indirect Branch Message	6	6	fixed	TCODE number = 4
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	1	8	variable	# sequential instructions executed since last taken branch
	1	32	variable	unique part of target address for taken branches/exceptions
Error Message	6	6	fixed	TCODE number = 8
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	5	5	fixed	error code

Table 9-2. Public TCODEs Supported (continued)

Message Name	Minimum Packet Size (bits)	Maximum Packet Size (bits)	Packet Type	Packet Description
Program Trace—Direct Branch Message w/ Sync	6	6	fixed	TCODE number = 11
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	1	8	variable	# sequential instructions executed since last taken branch
	1	32	variable	full target address (leading zeros truncated)
Program Trace—Indirect Branch Message w/ Sync	6	6	fixed	TCODE number = 12
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	1	8	variable	# sequential instructions executed since last taken branch
	1	32	variable	full target address (leading zeros truncated)
Watchpoint Message	6	6	fixed	TCODE number = 15
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	8	8	fixed	# indicating watchpoint source(s)
Resource Full Message	6	6	fixed	TCODE number = 27
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	4	4	fixed	resource code (Refer to Table 9-4)—indicates which resource is the cause of this message
	1	32	variable	branch / predicate instruction history (see Section)
Program Trace—Indirect Branch History Message	6	6	fixed	TCODE number = 28 (see Note below)
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	1	8	variable	# sequential instructions executed since last taken branch
	1	32	variable	unique part of target address for taken branches/exceptions
	1	32	variable	branch / predicate instruction history (see Section 9.8.1, “Branch Trace Messaging (BTM)”)
Program Trace—Indirect Branch History Message w/ Sync	6	6	fixed	TCODE number = 29 (see Note below)
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	1	8	variable	# sequential instructions executed since last taken branch
	1	32	variable	full target address (leading zero (0) truncated)
	1	32	variable	branch / predicate instruction history (see Section 9.8.1, “Branch Trace Messaging (BTM)”)
Program Trace—Program Correlation Message	6	6	fixed	TCODE number = 33
	4	4	fixed	source processor identifier (multiple Nexus configuration)
	4	4	fixed	event correlated w/ program flow (Refer to Table 9-5)
	1	8	variable	# sequential instructions executed since last taken branch
	1	32	variable	branch / predicate instruction history (see Section 9.8.1, “Branch Trace Messaging (BTM)”)

Table 9-3 shows the error code encodings used when reporting an error via the Nexus 2+ Error Message.

Table 9-3. Error Code Encoding (TCODE = 8)

Error Code	Description
00000	Ownership Trace overrun
00001	Program Trace overrun
00010	Reserved for Data Trace overrun (not in Nexus 2+)
00011	Read/write access error
00101	Invalid access opcode (Nexus Register unimplemented)
00110	Watchpoint overrun
00111	Program Trace and Ownership Trace overrun
01000	(Program Trace or Ownership Trace) and Watchpoint overrun
01001–10111	Reserved
11000	BTM lost due to collision w/ higher priority messages
11001–11111	Reserved

Table 9-4 shows the encodings used for resource codes for certain messages.

Table 9-4. RCODE values (TCODE = 27)

Resource Code	Description
0000	Program Trace Instruction counter reached 255 and was reset.
0001	Program Trace, Branch / Predicate Instruction History. This type of packet is terminated by a stop bit set to 1 after the last history bit.

Table 9-5 shows the event code encodings used for certain messages.

Table 9-5. Event Code Encoding (TCODE = 33)

Event Code	Description
0000	Entry into Debug Mode
0001	Entry into Low Power Mode (CPU only)
0010-0011	Reserved for future functionality
0100	Disabling Program Trace
0101–1111	Reserved for future functionality

Table 9-6 shows the data trace size encodings used for certain messages.

Table 9-6. Data Trace Size Encodings (TCODE = 5,6,13,14)

DTM Size Encoding	Transfer Size
000	Byte
001	Halfword (2 bytes)
010	Word (4 bytes)
011	Reserved
100	String (3 bytes)
101–111	Reserved

NOTE

Program Trace can be implemented using either Branch History/Predicate Instruction Messages, or traditional Direct/Indirect Branch Messages. The user can select between the two types of Program Trace. The advantages for each are discussed in [Section 9.8.1, “Branch Trace Messaging \(BTM\)”](#). If the Branch History method is selected, the shaded TCODES above are not messaged out.

9.4 Nexus 2+ Programmer’s Model

This section describes the Nexus 2+ programmers model. Nexus 2+ registers are accessed using the JTAG/OnCE port in compliance with IEEE 1149.1. See [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) for details on Nexus 2+ register access.

NOTE

Nexus 2+ registers and output signals are numbered using bit 0 as the least significant bit. This bit ordering is consistent with the ordering defined by IEEE-ISTO 5001.

details the register map for the Nexus 2+ module.

Table 9-7. Nexus 2+ Register Map

Nexus Register	Nexus Access Opcode	Read/Write	Read Address	Write Address
Client Select Control (CSC) ¹	0x1	R	0x02	—
Port Configuration Register (PCR) ¹	PCR_INDEX ²	R/W	—	—
Development Control1 (DC1)	0x2	R/W	0x04	0x05
Development Control2 (DC2)	0x3	R/W	0x06	0x07
Development Status (DS)	0x4	R	0x08	—
Read/Write Access Control/Status (RWCS)	0x7	R/W	0x0E	0x0F
Read/Write Access Address (RWA)	0x9	R/W	0x12	0x13
Read/Write Access Data (RWD)	0xA	R/W	0x14	0x15
Watchpoint Trigger (WT)	0xB	R/W	0x16	0x17
Reserved	0xC -> 0x3F	—	0x1A->0x7E	0x19->7F

¹ The CSC and PCR registers are shown in this table as part of the Nexus programmer's model. They are only present at the top level SoC Nexus controller in a multi-Nexus implementation, not in the Nexus 2+ module. The SoC's CSC Register is readable through Nexus, but the PCR is shown for reference only here.

² The "PCR_INDEX" is a parameter determined by the SoC.

9.4.1 Client Select Control (CSC)

The CSC Register determines which Nexus client is under development. This register is present at the top-level SOC Nexus 2+ controller to select one of multiple on-chip Nexus 2+ units.




Figure 9-2. Client Select Control Register

Table 9-8. Client Select Control Register Fields

CSC[7:4]	RES—Reserved for future Nexus Clients (read as 0)
CSC[3:0]	CSC—Client Select Control 0xX = Nexus client (SoC level)

9.4.2 Port Configuration Register (PCR)

The Port Configuration Register (PCR) controls the basic port functions for all Nexus modules in a multi-Nexus environment. This includes clock control and auxiliary port width. All bits in this register are writable only once after system reset.




Figure 9-3. Port Configuration Register

Table 9-9. Port Configuration Register Fields

PCR[31]	OPC	OPC—Output Port Mode Control 0 Reduced Port Mode configuration (min# <i>nex_mdo</i> [n:0] pins defined by SOC) 1 Full Port Mode configuration (max# <i>nex_mdo</i> [n:0] pins defined by SOC)
PCR[30]	—	Reserved for future functionality
PCR[29]	MCK_EN	MCK_EN—MCKO Clock Enable 0 <i>nex_mcko</i> is disabled 1 <i>nex_mcko</i> is enabled
PCR[28:26]	MCK_DIV	MCK_DIV—MCKO Clock Divide Ratio (see note below) 000 <i>nex_mcko</i> is 1x processor clock freq. 001 <i>nex_mcko</i> is 1/2x processor clock freq. 010 Reserved (default to 1/2x processor clock freq.) 011 <i>nex_mcko</i> is 1/4x processor clock freq. 100–110 Reserved (default to 1/2x processor clock freq.) 111 <i>nex_mcko</i> is 1/8x processor clock freq.
PCR[25:0]	—	Reserved for future functionality

NOTE

The CSC and PCR Registers exist in a separate module at the SoC level in a multi-Nexus environment. If the e200 Nexus 2+ module is the only Nexus module, these registers are not implemented and the e200 Nexus 2+ defined Development Control Register 1 (DC1) is used to control Nexus port functionality.

9.4.3 Development Control Register 1, 2 (DC1, DC2)

The Development Control Registers are used to control the basic development features of the Nexus 2+ module. Development Control Register 1 is shown in [Figure 9-4](#) and its fields are described in [Table 9-10](#).




Figure 9-4. Development Control Register 1

Table 9-10. Development Control Register 1 Fields

DC1[31]	OPC	OPC—Output Port Mode Control 0 Reduced Port Mode configuration (min# <i>nex_mdo</i> [n:0] pins defined by SOC) 1 Full Port Mode configuration (max# <i>nex_mdo</i> [n:0] pins defined by SOC)
DC1[30:29]	MCK_DIV	MCK_DIV—MCKO Clock Divide Ratio (see note below) 00 <i>nex_mcko</i> is 1x processor clock freq. 01 <i>nex_mcko</i> is 1/2x processor clock freq. 10 <i>nex_mcko</i> is 1/4x processor clock freq. 11 <i>nex_mcko</i> is 1/8x processor clock freq.
DC1[28:27]	EOC	EOC—EVTO Control 00 <i>nex_evto_b</i> upon occurrence of Watchpoints (configured in DC2) 01 <i>nex_evto_b</i> upon entry into Debug Mode 10 <i>nex_evto_b</i> upon Timestamping Event 11 Reserved
DC1[26]	—	Reserved for future functionality
DC1[25]	PTM	PTM—Program Trace Method 0 Program Trace uses traditional Branch Messages 1 Program Trace uses Branch History Messages
DC1[24]	WEN	WEN—Watchpoint Trace Enable 0 Watchpoint Messaging disabled 1 Watchpoint Messaging enabled
DC1[23:8]	—	Reserved for future functionality
DC1[7:5]	OVC	OVC—Overrun Control 000 Generate overrun messages 001–010 Reserved 011 Delay processor for BTM / OTM overruns 1XX Reserved
DC1[4:3]	EIC	EIC—EVTI Control 00 <i>nex_evti_b</i> is used for synchronization (Program Trace/ Data Trace) 01 <i>nex_evti_b</i> is used for Debug request 1X Reserved
DC1[2:0]	TM	TM—Trace Mode 000 No Trace 1XX Program Trace enabled X1X Reserved XX1 Ownership Trace enabled

NOTE

The Output Port Mode Control bit (OPC) and MCKO Clock Divide Ratio bits (MCK_DIV) MUST ONLY be modified during system reset or debug mode to insure correct output port and output clock functionality. It is also recommended that all other bits of the DC1 also only be modified in one of these two modes.

Development Control Register 2 is shown in [Figure 9-5](#) and its fields are described in [Table 9-11](#).




Figure 9-5. Development Control Register 2

Table 9-11. Development Control Register 2 Fields

DC2[31:24]	EWC	EWC—EVTO Watchpoint Configuration 00000000 No Watchpoints trigger <i>nex_evto_b</i> 1XXXXXXX Watchpoint #0 (IAC1 from Nexus1) triggers <i>nex_evto_b</i> X1XXXXXX Watchpoint #1 (IAC2 from Nexus1) triggers <i>nex_evto_b</i> XX1XXXXX Watchpoint #2 (IAC3 from Nexus1) triggers <i>nex_evto_b</i> XXX1XXXX Watchpoint #3 (IAC4 from Nexus1) triggers <i>nex_evto_b</i> XXXX1XXX Watchpoint #4 (DAC1 from Nexus1) triggers <i>nex_evto_b</i> XXXXX1XX Watchpoint #5 (DAC2 from Nexus1) triggers <i>nex_evto_b</i>
DC2[23:0]	—	Reserved for future functionality

NOTE

The EOC bits in DC1 must be programmed to trigger $\overline{\text{EVTO}}$ on Watchpoint occurrence for the EWC bits to have any effect.

9.4.4 Development Status Register (DS)

The Development Status Register is used to report system debug status. When Debug Mode is entered or exited, or an SOC or Zen defined Low Power Mode is entered (see Note below), a Debug Status Message is transmitted with DS[31:24]. The external tool can read this register at any time.





Figure 9-6. Development Status Register

Table 9-12. Development Status Register Fields

DS[31]	DBG	DBG—Zen CPU Debug Mode Status 0 CPU not in Debug mode 1 CPU in Debug mode (<i>jd_debug_b</i> signal asserted)
DS[30:28]	LPS	LPS—Zen System Low Power Mode Status 000 Normal (Run) mode XX1 DOZE mode (<i>p_doze</i> signal asserted) X1X NAP mode (<i>p_nap</i> signal asserted) 1XX SLEEP mode (<i>p_sleep</i> signal asserted)
DS[27:26]	LPC	LPC—Zen CPU Low Power Mode Status 00 Normal (Run) mode 01 CPU in Halted state (<i>p_halted</i> signal asserted) 10 CPU in Stopped state (<i>p_stopped</i> signal asserted) 11 CPU in Waiting state (<i>p_waiting</i> signal asserted)
DS[25]	CHK	CHK—Zen CPU Checkstop Status 0 CPU not in Checkstop state 1 CPU in Checkstop state (<i>p_chkstop</i> signal asserted)
DS[24:0]	—	Reserved for future functionality (read as 0)

9.4.5 Read/Write Access Control/Status (RWCS)

The Read Write Access Control/Status Register provides control for Read/Write Access. Read/Write access provides DMA-like access to memory-mapped resources on the AHB bus either while the processor is halted, or during runtime. The RWCS Register also provides Read/Write Access Status information per [Table 9-14](#).


Figure 9-7. Read/Write Access Control/Status Register

¹ ERR and DV are read-only

Table 9-13. Read/Write Access Control/Status Register Fields

RWCS[31]	AC	AC—Access Control 0 End access 1 Start access
RWCS[30]	RW	RW—Read/Write Select 0 Read access 1 Write access

Table 9-13. Read/Write Access Control/Status Register Fields (continued)

RWCS[29:27]	SZ	SZ—Word Size 000 8-bit (byte) 001 16-bit (half-word) 010 32-bit (word) 011 Reserved 100–111 Reserved (default to word)
RWCS[26:24]	MAP	MAP—MAP Select 000 Primary memory map 001–111 Reserved
RWCS[23:22]	PR	PR—Read/Write Access Priority 00 Lowest access priority 01 Reserved (default to lowest priority) 10 Reserved (default to lowest priority) 11 Highest access priority
RWCS[21:16]	—	RES—Reserved for future functionality
RWCS[15:2]	CNT	CNT—Access Control Count hhhh Number of accesses of word size SZ
RWCS[1]	ERR ¹	ERR—Read/Write Access Error (see Table 9-14)
RWCS[0]	DV ¹	DV—Read/Write Access Data Valid (see Table 9-14)

¹ ERR and DV are read-only

Table 9-14. Read/Write Access Status Bit Encoding

Read Action	Write Action	ERR	DV
Read Access has not completed	Write Access completed without error	0	0
Read Access error has occurred	Write Access error has occurred	1	0
Read Access completed without error	Write Access has not completed	0	1
Not Allowed	Not allowed	1	1

9.4.6 Read/Write Access Data (RWD)

The Read/Write Access Data Register (RWD) provides the data to/from AHB bus memory-mapped locations when initiating a read or a write access.




Figure 9-8. Read/Write Access Data Register

Read/Write accesses to the AHB bus require that the debug firmware properly retrieve/place the data in

the RWD. Table 9-15 shows the proper placement of data into the RWD.

Table 9-15. RWD data placement for Transfers

Transfer Size and byte offset	RWA(2:0)	RWCS[SZ]	RWD			
			31:24	23:16	15:8	7:0
Byte	x x x	0 0 0	—	—	—	X
Half	x x 0	0 0 1	—	—	X	X
Word	x 0 0	0 1 0	X	X	X	X

Note:

“X” indicates byte lanes with valid data

“—” indicates byte lanes that contain unused data.

Table 9-16 shows the mapping of RWD bytes to byte lanes of the AHB system bus read and write data buses.

Table 9-16. RWD byte lane data placement

Transfer Size and byte offset	RWA(2:0)	RWD			
		31:24	23:16	15:8	7:0
Byte @000	0 0 0	—	—	—	AHB[7:0]
Byte @001	0 0 1	—	—	—	AHB[15:8]
Byte @010	0 1 0	—	—	—	AHB[23:16]
Byte @011	0 1 1	—	—	—	AHB[31:24]
Half @000	0 0 0	—	—	AHB[15:8]	AHB[7:0]
Half @010	0 1 0	—	—	AHB[31:24]	AHB[23:16]
Word @000	0 0 0	AHB[31:24]	AHB[23:16]	AHB[15:8]	AHB[7:0]

Note:

“—” indicates byte lanes that contain unused data.

9.4.7 Read/Write Access Address (RWA)

The Read/Write Access Address Register provides the AHB system bus address to be accessed when initiating a read or a write access.




Figure 9-9. Read/Write Access Address Register

9.4.8 Watchpoint Trigger Register (WT)

The Watchpoint Trigger Register allows the watchpoints defined within the Zen Nexus1 logic to trigger actions. These watchpoints can control Program and/or Data Trace enable and disable. The WT bits can be used to produce an address related “window” for triggering Trace Messages.




Figure 9-10. Watchpoint Trigger Register

Table 9-17 details the Watchpoint Trigger register fields.

Table 9-17. Watchpoint Trigger Register Fields

WT[31:29]	PTS	PTS—Program Trace Start Control 000 Trigger disabled 001 Use Watchpoint #0 (IAC1 from Nexus1) 010 Use Watchpoint #1 (IAC2 from Nexus1) 011 Use Watchpoint #2 (IAC3 from Nexus1) 100 Use Watchpoint #3 (IAC4 from Nexus1) 101 Use Watchpoint #4 (DAC1 from Nexus1) 110 Use Watchpoint #5 (DAC2 from Nexus1) 111 Reserved
WT[28:26]	PTE	PTE—Program Trace End Control 000 Trigger disabled 001 Use Watchpoint #0 (IAC1 from Nexus1) 010 Use Watchpoint #1 (IAC2 from Nexus1) 011 Use Watchpoint #2 (IAC3 from Nexus1) 100 Use Watchpoint #3 (IAC4 from Nexus1) 101 Use Watchpoint #4 (DAC1 from Nexus1) 110 Use Watchpoint #5 (DAC2 from Nexus1) 111 Reserved
WT[25:0]	—	RES—Reserved for future functionality (read as 0)

NOTE

The WT bits ONLY control Program Trace if the TM bits within the Development Control Register 1 (DC1) have not already been set to enable Program Trace.

If the TM bits are set to disable Program Trace, then a value of 0 for PTS causes Instruction Trace to remain disabled regardless of the setting of PTE. Also, once triggered, (with the TM bits set to disable Program Trace) writing a value of 0 to PTS causes Instruction Trace to be disabled.

9.5 Nexus 2+ Register Access via JTAG/OnCE

Access to Nexus 2+ register resources is enabled by loading a single instruction (“NEXUS2-ACCESS”) into the JTAG Instruction Register (IR) (OnCE OCMD register). For the Nexus 2+ block, the OCMD value is 0b0001111100.

Once the “NEXUS2-ACCESS” instruction has been loaded, the JTAG/OnCE port allows tool/target communications with all Nexus 2+ registers according to the register map in [Table 9-7](#).

Reading/writing of a Nexus 2+ register then requires two (2) passes through the Data-Scan (DR) path of the JTAG state machine (see [Section 9.15, “IEEE 1149.1 \(JTAG\) RD/WR Sequences”](#)).

1. The first pass through the DR selects the Nexus 2+ register to be accessed by providing an index (see [Table 9-7](#)), and the direction (read/write). This is achieved by loading an 8-bit value into the JTAG Data Register (DR). This register has the following format:


(7bits)		(1 bit)
Nexus Register Index		R/W
RESET Value: 0x00		

Nexus Register Index:	Selected from values in Table 9-7
Read/Write (R/W):	0 = Read 1 = Write

2. The second pass through the DR then shifts the data in or out of the JTAG port, LSB first.
 - a) During a read access, data is latched from the selected Nexus register when the JTAG state machine passes through the “Capture-DR” state.
 - b) During a write access, data is latched into the selected Nexus register when the JTAG state machine passes through the “Update-DR” state.

9.6 Debug Status Messages

Debug Status Messages report low power mode and debug status. Debug Status Messages are enabled when Nexus 2+ is enabled. Entering/exiting Debug Mode as well as entering a Low Power Mode triggers a Debug Status Message, indicating the value of the most significant byte in the Development Status register. Debug status information is sent out in the following format:



Fixed length = 18 bits

Figure 9-11. Debug Status Message Format

9.7 Ownership Trace

This section details the ownership trace features of the Nexus 2+ module.

9.7.1 Overview

Ownership trace provides a macroscopic view, such as task flow reconstruction, when debugging software written in a high level (or object-oriented) language. It offers the highest level of abstraction for tracking operating system software execution. This is especially useful when the developer is not interested in debugging at lower levels.

9.7.2 Ownership Trace Messaging (OTM)

Ownership trace information is messaged via the auxiliary port using an Ownership Trace Message (OTM). Zen processors contain a Power Architecture BookE defined “Process ID” register within the CPU. It is updated by the operating system software to provide task/process ID information. The contents of this register are replicated on the pins of the processor and connected to Nexus. The Process ID register value can be accessed using the **mfspr/mtspr** instructions. Please refer to the Programmer’s Model section of the appropriate “Zen Implementation Definition” document for more details on the Process ID register.


NOTE

The CPU includes a Process ID register (PID0), thus the Nexus UBA functionality is not implemented.

There is one condition that causes an Ownership Trace Message, as follows:

1. When new information is updated in the Process ID register by the Zen processor, the data is latched within Nexus, and is messaged out via the auxiliary port, allowing development tools to trace ownership flow.

Ownership trace information is messaged out in the following format



Fixed length = 42 bits

Figure 9-12. Ownership Trace Message Format

9.7.3 OTM Error Messages

An Error Message occurs when a new message cannot be queued due to the message queue being full. The FIFO discards incoming messages until it has completely emptied the queue. Once emptied, an Error Message is queued. The error encoding indicates which type(s) of messages attempted to be queued while the FIFO was being emptied.

If only an OTM Message attempts to enter the queue while it is being emptied, the Error Message incorporates the OTM only error encoding (00000). If both OTM and either BTM or DTM messages attempt to enter the queue, the Error Message incorporates the OTM and (Program or Data) Trace error encoding (00111). If a Watchpoint also attempts to be queued while the FIFO is being emptied, then the Error Message incorporates error encoding (01000).

NOTE

The OVC bits within the DC1 Register can be set to delay the CPU in order to alleviate (but not eliminate) potential overrun situations.

Error information is messaged out in the following format (see [Table 9-3](#)).

(5 bits)	(4 bits)	(6 bits)
Error Code (00000 / 00111 / 01000)	Src. Proc.	TCODE (001000)

Fixed length = 15 bits

Figure 9-13. Error Message Format

9.7.4 OTM Flow

Ownership Trace Messages are generated when the operating system writes to the Zen Process ID register.

The following flow describes the OTM process.

1. The Process ID register is a system control register. It is internal to the Zen processor and can be accessed by using PPC instructions. The contents of this register are replicated on the pins of the processor and connected to Nexus.
2. Writes to the Zen internal Process ID register pulse a write signal to Nexus. The data value written into the Process ID register is latched and formed into the Ownership Trace Message that is queued to be transmitted.
3. Process ID register reads do not cause Ownership Trace Messages to be transmitted by the Nexus 2+ module.

9.8 Program Trace

This section details the program trace mechanism supported by Nexus 2+ for the e200 processor. Program trace is implemented via Branch Trace Messaging (BTM) as per the IEEE-ISTO 5001-2003 definition. Branch Trace Messaging for Zen processors is accomplished by snooping the Zen virtual address bus (between the CPU and MMU), attribute signals, and CPU Status (*p_pstat*[0:5]).

9.8.1 Branch Trace Messaging (BTM)

Traditional Branch Trace Messaging facilitates program trace by providing the following types of information:

- Messaging for taken direct branches includes how many sequential instructions were executed since the last taken branch or exception, including the taken direct branch. Branch instructions are included in the count of sequential instructions.
- Messaging for taken indirect branches and exceptions includes how many sequential instructions were executed since the last taken branch or exception and the unique portion of the branch target address or exception vector address. Branch instructions are included in the count of sequential instructions. For taken indirect branches which trigger generation of a message, the branch is also included in the count.

Branch History Messaging facilitates program trace by providing the following information.

- Messaging for taken indirect branches and exceptions includes a) how many sequential instructions (I-CNT) were executed since the last predicate instruction, taken/not taken direct branch, taken/not-taken indirect branch, or exception, b) the unique portion of the branch target address or exception vector address, and c) a branch/predicate instruction history field. Each bit in the history field represents a direct branch or predicated instruction where a value of one (1) indicates taken, and a value of zero (0) indicates not taken. Not-taken indirect branches generate a history bit with a value of zero (0). Instructions that generate history bits are not included in instruction counts. For taken indirect branches which trigger generation of this message type, the branch is included in the count, but not in the history field

9.8.1.1 Zen Indirect Branch Message Instructions

Table 9-18 shows the types of instructions and events which cause Indirect Branch Messages or Branch History Messages to be encoded.

Table 9-18. Indirect Branch Message Sources

Source of Indirect Branch Message	Instructions / Detail
Taken branch relative to a register value	se_bctr, se_bctrl, se_blr, se_blrl
System Call/Trap exceptions taken	sc, se_sc, tw
Return from interrupts / exceptions	se_rfi, se_rfci, se_rfdi
Exit from reset with Program Trace Enabled	Indirect branch with Sync, target address is initial instruction, count=1

9.8.1.2 Zen Direct Branch Message Instructions

Table 9-19 shows the types of instructions that cause Direct Branch Messages or toggle a bit in the instruction history buffer to be messaged out in a Resource Full Message or Branch History Message.

Table 9-19. Direct Branch Message Sources

Source of Direct Branch Message	Instructions
Taken direct branch instructions Instruction Synchronize	se_b, se_bc, se_bl, e_b, e_bc, e_bl, e_bcl, se_isync

9.8.1.3 BTM using Branch History Messages

Traditional BTM Messaging can accurately track the number of sequential instructions between branches, but cannot accurately indicate which instructions were conditionally executed, and which were not.

Branch History Messaging solves this problem by providing a predicated instruction history field in each Indirect Branch Message. Each bit in the history represents a predicated instruction or direct branch, or a not-taken indirect branch. A value of one (1) indicates the conditional instruction was executed or the direct branch was taken. A value of zero (0) indicates the conditional instruction was not executed or the branch was not taken.

Branch History Messages solve predicated instruction tracking and save bandwidth because only indirect branches cause messages to be queued.

9.8.1.4 BTM using Traditional Program Trace Messages

Based on the PTM bit in the DC1 Register (DC1[25]), Program Tracing can utilize either Branch History Messages (DC1[25]=1'b1) or traditional Direct/Indirect Branch Messages (DC1[25]=1'b0).

Branch History saves bandwidth and keep consistency between methods of Program Trace, yet may lose temporal order between BTM messages and other types of messages. Because direct branches are not messaged, but are instead included in the history field of the Indirect Branch History Message, other types of messages may enter the FIFO between Branch History Messages. The development tool cannot determine the ordering of “events” that occurred with respect to direct branches simply by the order in which messages are sent out.

Traditional BTM messages maintain their temporal ordering because each event that can cause a message to be queued enters the FIFO in the order it occurred and is messaged out maintaining that order.

9.8.2 BTM Message Formats

The Nexus 2+ block supports three types of traditional BTM Messages—Direct, Indirect, and Synchronization Messages. It supports two types of branch history BTM Messages—Indirect Branch History, and Indirect Branch History with Synchronization Messages. Program Correlation, Resource Full, and Error Messages are also supported.

9.8.2.1 Indirect Branch Messages (History)

Indirect branches include all taken branches whose destination is determined at run time, interrupts, and exceptions. If DC1[25] is set, indirect branch information is messaged out in the following format:

(1-32 bits)	(1-32 bits)	(1-8 bits)	(4 bits)	(6 bits)
Branch History	Relative Address	Sequence Count	Src. Proc.	TCODE (011100)

Max length = 82 bits; Min length = 13 bits

Figure 9-14. Indirect Branch Message (History) Format

9.8.2.2 Indirect Branch Messages (Traditional)

If DC1[25] is cleared, indirect branch information is messaged out in the following format:

(1-32 bits)	(1-8 bits)	(4 bits)	(6 bits)
Relative Address	Sequence Count	Src. Proc.	TCODE (000100)

Max length = 50 bits; Min length = 12 bits

Figure 9-15. Indirect Branch Message Format

9.8.2.3 Direct Branch Messages (Traditional)

Direct branches (conditional or unconditional) are all taken branches whose destination is fixed in the instruction opcode. Direct branch information is messaged out in the following format

(1-8 bits)	(4 bits)	(6 bits)
Sequence Count	Src. Proc.	TCODE (000011)

Max length = 18 bits; Min length = 11bits

Figure 9-16. Direct Branch Message Format

NOTE

When DC1[25] is set, Direct Branch Messages are not transmitted. Instead, each direct branch, not-taken indirect branch, or predicated instruction is recorded in the history buffer.

9.8.2.4 Resource Full Messages

The Resource Full Message is used in conjunction with Branch Trace and Branch History Messages. The Resource Full Message is generated when either the internal branch/predicate history buffer is full, or if the BTM Instruction sequence counter (I-CNT) overflows. If synchronization is needed at the time this

message is generated, the synchronization is delayed until the next Branch Trace Message that is not a Resource Full Message.

For history buffer overflow, the Resource Full Message transmits a Resource Code (RCODE) of 0b0001 and the current contents of the history buffer, including the stop bit, are transmitted in the Resource Data (RDATA) field. This history information can be concatenated by the development tool with the branch/predicate history information from subsequent messages to obtain the complete branch/predicate history between indirect changes of flow.

For instruction counter overflow, the Resource Full Message transmits an RCODE of 0b0000 and a value of 0xFF is transmitted in the RDATA field, indicating that 255 sequential instructions have been executed since the last change of flow or, if program trace is in history mode, since the last instruction that recorded history information.




Figure 9-17. Resource Full Message Format

Table 9-20 shows the RCODE encodings and RDATA information used for Resource Full messages.

Table 9-20. RCODE Encoding

RCODE	Description	RDATA field
0000	Program Trace Instruction counter reached 255 and was reset.	0xFF
0001	Program Trace, Branch / Predicate Instruction History full.	Branch History. This type of packet is terminated by a stop bit set to 1 after the last history bit.

9.8.2.5 Program Correlation Messages

Program Correlation Messages (PCMs) are used to correlate events to the program flow that may not be associated with the instruction stream. The following events result in a PCM when program trace is enabled:

- When the CPU enters debug mode, a PCM is generated. The instruction count and history information provided by the PCM can be used to determine the last sequence of instructions executed prior to debug mode entry.
- When the CPU enters a low power mode in which instructions are no longer executed, a PCM is generated. The instruction count and history information provided by the PCM can be used to determine the last sequence of instructions executed prior to low power mode entry.
- Whenever program trace is disabled by any means, a PCM is generated. The instruction count and history information provided by the PCM can be used to determine the last sequence of instructions executed prior to disabling program trace.

Refer to [Table 9-5](#) for the event codes that are supported in this implementation. Program Correlation is messaged out in the following format:

(1-32 bits)	(1-8 bits)	(4 bits)	(4 bits)	(6 bits)
Branch History	Sequence Count	ECODE	Src. Proc.	TCODE (100001)

Max length = 54 bits; Min length = 16 bits

Figure 9-18. Program Correlation Message Format

9.8.2.6 BTM Overflow Error Messages

An Error Message occurs when a new message cannot be queued due to the message queue being full. The FIFO discards incoming messages until it has completely emptied the queue. Once emptied, an Error Message is queued. The error encoding indicates which type(s) of messages attempted to be queued while the FIFO was being emptied.

If only a Program Trace Message attempts to enter the queue while it is being emptied, the Error Message incorporates the Program Trace only error encoding (00001). If both OTM and Program Trace Messages attempt to enter the queue, the Error Message incorporates the OTM and Program Trace error encoding (00111). If a Watchpoint also attempts to be queued while the FIFO is being emptied, then the Error Message incorporates error encoding (01000).

NOTE

The OVC bits within the DC1 Register can be set to delay the CPU in order to alleviate (but not eliminate) potential overrun situations.

Error information is messaged out in the following format:

(5 bits)	(4 bits)	(6 bits)
Error Code (00001 / 00111 / 01000)	Src. Proc.	TCODE (001000)

Fixed length = 15 bits

Figure 9-19. Error Message Format

9.8.2.7 Program Trace Synchronization Messages

A Program Trace Direct/Indirect Branch with Sync Message is messaged via the auxiliary port (provided Program Trace is enabled) for the following conditions (see [Table 9-21](#)):

- Exit from reset with program trace already enabled
- Initial Program Trace Message upon the first direct/indirect branch after exit from system reset or whenever program trace is enabled.
- Upon direct/indirect branch after returning from a CPU Low Power state.
- Upon direct/indirect branch after returning from Debug Mode.
- Upon direct/indirect branch after occurrence of queue overrun (can be caused by any trace message), provided Program Trace is enabled.

- Upon direct/indirect branch after the periodic program trace counter has expired indicating 255 *without-sync* Program Trace Messages have occurred since the last *with-sync* message occurred.
- Upon direct/indirect branch after assertion of the Event In (*nex_evti_b*) pin if the EIC bits within the DC1 Register have enabled this feature.
- Upon direct/indirect branch after the sequential instruction counter has expired indicating 255 instructions have occurred since the last change of flow.
- Upon direct/indirect branch after a BTM Message was lost due to an attempted access to a secure memory location (for SOCs with security).
- Upon direct/indirect branch after a BTM Message was lost due to a collision entering the FIFO between the BTM Message and both a Watchpoint Message and an Ownership Trace Message.
- Upon the first direct/indirect branch message after an execution mode switch into or out of a sequence of VLE instructions.

The format for Program Trace Direct/Indirect Branch with Sync Messages is as follows:

(1-32 bits)	(1-8 bits)	(4 bits)	(6 bits)
Full Target Address	Sequence Count	Source Proc.	TCODE (001011 or 001100)

Max length = 50 bits; Min length = 12 bits

Figure 9-20. Direct/Indirect Branch w/ Sync. Message Format

The formats for Program Trace Direct/Indirect Branch with Sync. Messages and Indirect Branch History with Sync. Messages are as follows

(1-32 bits)	(1-32 bits)	(1-8 bits)	(4 bits)	(6 bits)
Branch History	Full Target Address	Sequence Count	Source Proc.	TCODE (011101)

Max length = 82 bits; Min length = 13 bits

Figure 9-21. Indirect Branch History w/ Sync. Message Format

Exception conditions that result in Program Trace Synchronization are summarized in [Table 9-21](#).

Table 9-21. Program Trace Exception Summary

Exception Condition	Exception Handling
System Reset Negation	At the negation of JTAG reset (<i>j_trst_b</i>), queue pointers, counters, state machines, and registers within the Nexus 2+ module are reset. Upon exiting system reset, if Program Trace is already enabled), a Program Trace Message is sent as an Indirect Branch w/ Sync. Message.
Program Trace Enabled	The first Program Trace Message (after Program Trace has been enabled) is a synchronization message.
Exit from Low Power/Debug	Upon exit from a Low Power mode or Debug mode the next direct/indirect branch is converted to a Direct/Indirect Branch with Sync. Message.

Table 9-21. Program Trace Exception Summary (continued)

Exception Condition	Exception Handling
Queue Overrun	An Error Message occurs when a new message cannot be queued due to the message queue being full. The FIFO discards messages until it has completely emptied the queue. Once emptied, an Error Message is queued. The error encoding indicates which type(s) of messages attempted to be queued while the FIFO was being emptied. The next BTM message in the queue is a Direct/Indirect Branch w/ Sync. Message.
Periodic Program Trace Sync.	A forced synchronization occurs periodically after 255 non-sync Program Trace Messages have been queued. A Direct/Indirect Branch w/ Sync. Message is queued. The periodic program trace message counter then resets.
Event In	If the Nexus module is enabled, a <i>nex_evti_b</i> assertion initiates a Direct/Indirect Branch w/ Sync. Message upon the next direct/indirect branch (if Program Trace is enabled and the EIC bits of the DC1 Register have enabled this feature).
Sequential Instruction Count Overflow	After the sequential instruction counter reaches its maximum count (up to 255 sequential instructions may be executed), a forced synchronization occurs. The sequential counter then resets. A Program Trace Direct/Indirect Branch w/ Sync. Message is queued upon execution of the next branch. A Resource Full Message is Queued on the overflow event. If a branch instruction is the 255th instruction to occur, and causes a Program Trace message to be queued, then no Resource Full Message is queued, and the w/Sync message is queued for the <i>next</i> Program Trace Direct/Indirect Branch Message.
Attempted Access to Secure Memory	For SOCs that implement security, any attempted branch to secure memory locations temporarily disables Program Trace and cause the corresponding BTM to be lost. The following direct/indirect branch queues a Direct/Indirect Branch w/ Sync. Message. The count value within this message is inaccurate because the re-enable of Program Trace is not necessarily aligned on an instruction boundary.
Collision Priority	All Messages have the following priority: WPM -> OTM -> BTM. A BTM Message which attempts to enter the queue at the same time as a Watchpoint Message and Ownership Trace Message is lost. An Error Message is sent indicating the BTM was lost. The following direct/indirect branch queues a Direct/Indirect Branch w/ Sync. Message. The count value within this message reflects the number of sequential instructions executed after the last successful BTM Message was generated. This count includes the branch that did not generate a message due to the collision.

9.8.3 BTM Operation

9.8.3.1 Enabling Program Trace

Branch Trace Messaging can be enabled in one of two ways:

- Setting the TM field of the DC1 Register to enable Program Trace (DC1[2]).
- Using the PTS field of the WT Register to enable Program Trace on Watchpoint hits (Zen watchpoints are configured within the CPU).

9.8.3.2 Relative Addressing

The relative address feature is compliant with the IEEE-ISTO 5001-2003 recommendations, and is designed to reduce the number of bits transmitted for addresses of Indirect Branch Messages.

The address transmitted is relative to the target address of the instruction which triggered the previous Indirect Branch (or Sync) Message. It is generated by XORing the new address with the previous address,

and then using only the results up to the most significant ‘1’ in the result. To recreate this address, an XOR of the (most-significant 0-padded) message address with the previously decoded address gives the current address.

Previous Address (A1) = 0x0003FC01, New Address (A2) = 0x0003F365




Figure 9-22. Relative Address Generation and Re-creation

9.8.3.3 Execution Mode Indication

In order for a development tool to properly interpret instruction count and history information, it must be aware of the execution mode context of that information. VLE instructions are interpreted differently from non-VLE instructions.

Program trace messages provide the execution mode status in the least significant bit of the reconstructed address field. A value of ‘0’ indicates that preceding instruction count and history information should be interpreted in a non-VLE context. A value of ‘1’ indicates that the preceding instruction count and history information should be interpreted in a VLE context. Note that when a branch results in an execution mode switch, the program trace message resulting from that branch indicates the previous execution state. The new state is not signaled until the next program trace message.

In some cases, a Program Correlation Message is generated to indicate execution mode status. Refer to [Section 9.8.2.5, “Program Correlation Messages,”](#) for more information on these cases.

9.8.3.4 Branch/Predicate Instruction History (HIST)

If DC[25] is set, BTM messaging uses the Branch History format. The branch history (HIST) packet in these messages provides a history of branch execution used for reconstructing the program flow. This packet is implemented as a left-shifting shift register. The register is always pre-loaded with a value of one

(1). This bit acts as a stop bit so that the development tools can determine which bit is the end of the history information. The pre-loaded bit itself is not part of the history, but is transmitted with the packet.

A value of one (1) is shifted into the history buffer on a taken direct branch (conditional or unconditional) and on any instruction whose predicate condition evaluated as true. A value of zero (0) is shifted into the history buffer on any instruction whose predicate condition evaluated as false as well as on branches not taken. This includes indirect as well as direct branches not taken.

9.8.3.5 Sequential Instruction Count (I-CNT)

The I-CNT packet is present in all BTM Messages. For traditional Branch Messages, I-CNT represents the number of sequential instructions including non-taken branches since the last Direct/Indirect Branch Messages. Branch instructions which trigger message generation are included in the I-CNT.

For Branch History Messages, I-CNT represents the number of instructions executed since the last taken/non-taken direct branch, predicate instruction, last taken/not-taken indirect branch, or exception. Branch instructions which trigger message generation are included in the I-CNT. Instructions which generate history bits are not included in the I-CNT.

The sequential instruction counter overflows after its value reaches 255 and is reset to 0. In addition, the next BTM Message (corresponding to the 256th or later instruction) is converted to a synchronization type message.

9.8.3.6 Program Trace Queueing

Nexus 2+ implements a programmable depth queue (32 minimum entry recommended) for queuing all messages. Messages that enter the queue are transmitted via the auxiliary pins in the order in which they are queued.

NOTE

If multiple trace messages need to be queued at the same time, Watchpoint Messages have the highest priority (WPM -> OTM -> BTM -> DTM). Up to two messages may be simultaneously queued.

9.8.4 Program Trace Timing Diagrams (2 MDO/1 MSEO configuration)




Figure 9-23. Program Trace—Indirect Branch Message (Traditional)




Figure 9-24. Program Trace—Indirect Branch Message (History)




Figure 9-25. Program Trace—Direct Branch (Traditional) and Error Messages




Figure 9-26. Program Trace—Indirect Branch w/ Sync. Message

9.9 Watchpoint Support

This section details the Watchpoint features of the Nexus 2+ module.

9.9.1 Overview

The Nexus 2+ module provides Watchpoint Messaging via the auxiliary pins, as defined by IEEE-ISTO 5001-2003.

Nexus 2+ is not compliant with Class4 Breakpoint/Watchpoint requirements defined in the standard. The Breakpoint/Watchpoint Control Register is not implemented.

9.9.2 Watchpoint Messaging

Enabling Watchpoint Messaging is done by setting the Watchpoint Enable bit in the DC1 Register. Setting the individual Watchpoint sources is supported through the Zen Nexus1 module. The Zen Nexus1 module is capable of setting multiple address and/or data watchpoints. Please refer to the Debug chapter for details on Watchpoint initialization.

When these watchpoints occur, a watchpoint event signal from the Nexus1 module causes a message to be sent to the queue to be messaged out. This message includes the watchpoint number indicating which watchpoint caused the message.

The occurrence of any of the e200 defined watchpoints can be programmed to assert the Event Out (*nex_evto_b*) pin for one (1) period of the output clock (*nex_mcko*) (see [Table 9-24](#) for details on *nex_evto_b*).

Watchpoint information is messaged out in the following format.




Figure 9-27. Watchpoint Message Format.

Table 9-22. Watchpoint Source Encoding

Watchpoint Source (8 Bits)	Watchpoint Description
00000001	Zen Watchpoint #0 (IAC1 from Nexus1)
00000010	Zen Watchpoint #1 (IAC2 from Nexus1)
00000100	Zen Watchpoint #2 (IAC3 from Nexus1)
00001000	Zen Watchpoint #3 (IAC4 from Nexus1)
00010000	Zen Watchpoint #4 (DAC1 from Nexus1)
00100000	Zen Watchpoint #5 (DAC2 from Nexus1)

9.9.3 Watchpoint Error Message

An Error Message occurs when a new message cannot be queued due to the message queue being full. The FIFO discards messages until it has completely emptied the queue. Once emptied, an Error Message is queued. The error encoding indicates which type(s) of messages attempted to be queued while the FIFO was being emptied.

If only a Watchpoint Message attempts to enter the queue while it is being emptied, the Error Message incorporates the Watchpoint only error encoding (00110). If an OTM and/or Program Trace and/or Data Trace Message also attempts to enter the queue while it is being emptied, the Error Message incorporates error encoding (01000).

NOTE

The OVC bits within the DC1 Register can be set to delay the CPU in order to alleviate (but not eliminate) potential overrun situations.

Error information is messaged out in the following format (see [Table 9-3](#)).

(5 bits)	(4 bits)	(6 bits)
Error Code (00110 / 01000)	Src. Proc.	TCODE (001000)

Fixed length = 15 bits

Figure 9-28. Error Message Format

9.9.4 Watchpoint Timing Diagram (2 MDO/1 MSEO Configuration)




Figure 9-29. Watchpoint Message and Watchpoint Error Message

9.10 Nexus 2+ Read/Write Access to Memory-Mapped Resources

The Read/Write access feature allows access to memory-mapped resources via the JTAG/OnCE port. The Read/Write mechanism supports single as well as block reads and writes to Zen bus resources.

The Nexus 2+ module is capable of accessing resources on the Zen system bus (AHB), with multiple configurable priority levels. Memory-mapped registers and other non-cached memory can be accessed via the standard memory map settings.

All accesses are setup and initiated by the Read/Write Access Control/Status Register (RWCS), as well as the Read/Write Access Address (RWA) and Read/Write Access Data Registers (RWD).

Using the Read/Write Access Registers (RWCS/RWA/RWD), memory mapped Zen AHB resources can be accessed through Nexus 2+. The following subsections describe the steps which are required to access memory-mapped resources.

NOTE

Read/Write Access can only access memory mapped resources when system reset is de-asserted and clocks are running.

Misaligned accesses are NOT supported in the e200 Nexus 2+ module.

9.10.1 Single Write Access

1. Initialize the Read/Write Access Address Register (RWA) through the access method outlined in [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) using the Nexus Register Index of 0x9 (see [Table 9-7](#)). Configure as follows:
 - Write Address -> 32h'xxxxxxxx (write address)
2. Initialize the Read/Write Access Control/Status Register (RWCS) through the access method outlined in [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) using the Nexus Register Index of 0x7 (see [Table 9-7](#)). Configure the bits as follows:
 - Access Control (AC) -> 1b'1 (to indicate start access)
 - Map Select (MAP) -> 3b'000 (primary memory map)
 - Access Priority (PR) -> 2b'00 (lowest priority)
 - Read/Write (RW) -> 1b'1 (write access)
 - Word Size (SZ) -> 3b'0xx (32-bit, 16-bit, 8-bit)
 - Access Count (CNT) -> 14h'0000 or 14h'0001(single access)

NOTE

Access Count (CNT) of 14'h0000 or 14'h0001 performs a single access.

3. Initialize the Read/Write Access Data Register (RWD) through the access method outlined in [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) using the Nexus Register Index of 0xA (see [Table 9-7](#)). Configure as follows:
 - Write Data -> 32h'xxxxxxxx (write data)

4. The Nexus block then arbitrates for the AHB bus and transfer the data value from the data buffer RWD Register to the memory mapped address in the Read/Write Access Address Register (RWA). When the access has completed without error (ERR=1'b0), Nexus asserts the *nex_rdy_b* pin (see [Table 9-24](#) for detail on *nex_rdy_b*) and clears the DV bit in the RWCS Register. This indicates that the device is ready for the next access.

NOTE

Only the *nex_rdy_b* pin as well as the DV and ERR bits within the RWCS provide Read/Write Access status to the external development tool.

9.10.2 Block Write Access

1. For a block write access, follow Steps 1, 2, and 3 outlined in [Section 9.10.1, “Single Write Access,”](#) to initialize the registers, but using a value greater than one (14'h0001) for the CNT field in the RWCS Register.
2. The Nexus block then arbitrates for the AHB system bus and transfer the first data value from the RWD Register to the memory mapped address in the Read/Write Access Address Register (RWA). When the transfer has completed without error (ERR=1'b0), the address from the RWA Register is incremented to the next word size (specified in the SZ field) and the number from the CNT field is decremented. Nexus then asserts the *nex_rdy_b* pin. This indicates that the device is ready for the next access.
3. Repeat Step 3 in [Section 9.10.1, “Single Write Access,”](#) until the internal CNT value is zero (0). When this occurs, the DV bit within the RWCS is cleared to indicate the end of the block write access.

NOTE

The actual RWA value as well as the CNT field within the RWCS are not changed when executing a block write access. The original values can be read by the external development tool at any time.

9.10.3 Single Read Access

1. Initialize the Read/Write Access Address Register (RWA) through the access method outlined in [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) using the Nexus Register Index of 0x9 (see [Table 9-7](#)). Configure as follows:
 - Read Address -> 32h'xxxxxxxx (read address)
2. Initialize the Read/Write Access Control/Status Register (RWCS) through the access method outlined in [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) using the Nexus Register Index of 0x7 (see [Table 9-7](#)). Configure the bits as follows:
 - Access Control (AC) -> 1b'1 (to indicate start access)
 - Map Select (MAP) -> 3b'000 (primary memory map)
 - Access Priority (PR) -> 2b'00 (lowest priority)
 - Read/Write (RW) -> 1b'0 (read access)
 - Word Size (SZ) -> 3b'0xx (32-bit, 16-bit, 8-bit)

— Access Count (CNT) -> 14h'0000 or 14h'0001(single access)

NOTE

Access Count (CNT) of 14'h0000 or 14'h0001 performs a single access.

3. The Nexus block then arbitrates for the AHB system bus and the read data is transferred from the AHB to the RWD Register. When the transfer is completed without error (ERR=1'b0), Nexus asserts the *nex_rdy_b* pin (see [Table 9-24](#) for detail on *nex_rdy_b*) and sets the DV bit in the RWCS Register. This indicates that the device is ready for the next access.
4. The data can then be read from the Read/Write Access Data Register (RWD) through the access method outlined in [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) using the Nexus Register Index of 0xA (see [Table 9-7](#)).

NOTE

Only the *nex_rdy_b* pin as well as the DV and ERR bits within the RWCS provide Read/Write Access status to the external development tool.

9.10.4 Block Read Access

1. For a block read access, follow Steps 1 and 2 outlined in [Section 9.10.3, “Single Read Access.”](#) to initialize the registers, but using a value greater than one (14'h0001) for the CNT field in the RWCS Register.
2. The Nexus block then arbitrates for the AHB system bus and the read data is transferred from the AHB to the RWD Register. When the transfer has completed without error (ERR=1'b0), the address from the RWA Register is incremented to the next word size (specified in the SZ field) and the number from the CNT field is decremented. Nexus then asserts the *nex_rdy_b* pin. This indicates that the device is ready for the next access.
3. The data can then be read from the Read/Write Access Data Register (RWD) through the access method outlined in [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) using the Nexus Register Index of 0xA (see [Table 9-7](#)).
4. Repeat Steps 3 and 4 in [Section 9.10.3, “Single Read Access,”](#) until the CNT value is zero (0). When this occurs, the DV bit within the RWCS is set to indicate the end of the block read access.

NOTE

The data values must be shifted out 32-bits at a time LSB first.

NOTE

The actual RWA value as well as the CNT field within the RWCS are not changed when executing a block read access. The original values can be read by the external development tool at any time.

9.10.5 Error Handling

The Nexus 2+ module handles various error conditions as follows:

9.10.5.1 Bus Read/Write Error

All address and data errors that occur on read/write accesses to the Zen AHB system bus return a transfer error encoding on the *p_hresp*[1:0] signals. If this occurs:

1. The access is terminated without re-trying (AC bit is cleared)
2. The ERR bit in the RWCS Register is set
3. The Error Message is sent (TCODE = 8) indicating Read/Write Error

9.10.5.2 Access Termination

The following cases are defined for sequences of the Read/Write protocol that differ from those described in the above sections.

1. If the AC bit in the RWCS Register is set to start Read/Write accesses and invalid values are loaded into the RWD and/or RWA, then an AHB access error may occur. This is handled as described above.
2. If a block access is in progress (all cycles not completed), and the RWCS Register is written, then the original block access is terminated at the boundary of the nearest completed access.
 - a) If the RWCS is written with the AC bit set, the next Read/Write access begins and the RWD can be written to/read from.
 - b) If the RWCS is written with the AC bit cleared, the Read/Write access is terminated at the nearest completed access. This method can be used to break (early terminate) block accesses.

9.10.6 Read/Write Access Error Message

The Read/Write Access Error Message is sent out when an AHB system bus access error (read or write) has occurred.

Error information is messaged out in the following format:

(5 bits)	(4 bits)	(6 bits)
Error Code (00011)	Src. Proc.	TCODE (001000)

Fixed length = 15 bits

Figure 9-30. Error Message Format

9.11 Nexus 2+ Pin Interface

This section details information regarding the Nexus 2+ pins and pin protocol.

The Nexus 2+ pin interface provides the function of transmitting messages from the messages queues to the external tools. It is also responsible for handshaking with the message queues.

9.11.1 Pins Implemented

The Nexus 2+ module implements one (1) *nex_evti_b* and one (1) *nex_mseo_b* or two (2) *nex_mseo_b*[1:0]. It also implements a configurable number of *nex_mdo*[n:0] pins, (1) *nex_rdy_b* pin, (1) *nex_evto_b* pin, and one (1) clock output pin (*nex_mcko*). The output pins are synchronized to the Nexus 2+ output clock (*nex_mcko*).

All Nexus 2+ input functionality is controlled through the JTAG/OnCE port in compliance with IEEE 1149.1 (see [Section 9.5, “Nexus 2+ Register Access via JTAG/OnCE,”](#) for details). The JTAG pins are incorporated as I/O to the Zen processor, and are further described in [Section 8.4.2, “JTAG/OnCE Pins.”](#)

Table 9-23. JTAG Pins for Nexus 2+

JTAG Pins	Input/Output	Description of JTAG Pins (included in Zen Nexus 1)
<i>j_tdo</i>	O	The Test Data Output (<i>j_tdo</i>) pin is the serial output for test instructions and data. <i>j_tdo</i> is three-stateable and is actively driven in the “Shift-IR” and “Shift-DR” controller states. <i>j_tdo</i> changes on the falling edge of <i>j_tclk</i> .
<i>j_tdi</i>	I	The Test Data Input (<i>j_tdi</i>) pin receives serial test instruction and data. TDI is sampled on the rising edge of <i>j_tclk</i> .
<i>j_tms</i>	I	The Test Mode Select (<i>j_tms</i>) input pin is used to sequence the OnCE controller state machine. <i>j_tms</i> is sampled on the rising edge of <i>j_tclk</i> .
<i>j_tclk</i>	I	The Test Clock (<i>j_tclk</i>) input pin is used to synchronize the test logic, and control register access through the JTAG/OnCE port.
<i>j_trst_b</i>	I	The Test Reset (<i>j_trst_b</i>) input pin is used to asynchronously initialize the JTAG/OnCE controller.

The auxiliary pins are used to send and receive messages and are described in [Table 9-24](#).

Table 9-24. Nexus 2+ Auxiliary Pins

Auxiliary Pins	Input/Output	Description of Auxiliary Pins
<i>nex_mcko</i>	O	Message Clock Out (<i>nex_mcko</i>) is a free running output clock to development tools for timing of <i>nex_mdo</i> [n:0] and <i>nex_mseo_b</i> [1:0] pin functions. <i>nex_mcko</i> is programmable through the DC1 Register.
<i>nex_mdo</i> [n:0]	O	Message Data Out (<i>nex_mdo</i> [n:0]) are output pins used for OTM, BTM, and DTM. External latching of <i>nex_mdo</i> [n:0] shall occur on the rising edge of the Nexus2+ clock (<i>nex_mcko</i>).
<i>nex_mseo_b</i> [1:0]	O	Message Start/End Out (<i>nex_mseo_b</i> [1:0]) are output pins which indicate when a message on the <i>nex_mdo</i> [n:0] pins has started, when a variable length packet has ended, and when the message has ended. External latching of <i>nex_mseo_b</i> [1:0] shall occur on the rising edge of the Nexus2+ clock (<i>nex_mcko</i>). One or two pin MSEO functionality is determined at integration time per SOC implementation
<i>nex_rdy_b</i>	O	Ready (<i>nex_rdy_b</i>) is an output pin used to indicate to the external tool that the Nexus block is ready for the next Read/Write Access. If Nexus2+ is enabled, this signal is asserted upon successful (without error) completion of an AHB system bus transfer (Nexus read or write) and is held asserted until the JTAG/OnCE state machine reaches the "Capture_DR" state. Upon exit from system reset or if Nexus2+ is disabled, <i>nex_rdy_b</i> remains de-asserted
<i>nex_evto_b</i>	O	Event Out (<i>nex_evto_b</i>) is an output which, when asserted, indicates one of two events has occurred based on the EOC bits in the DC1 Register. <i>nex_evto_b</i> is held asserted for one (1) cycle of <i>nex_mcko</i> : - One (or more) watchpoints has occurred (from Nexus1) and EOC = 2'b00 - Debug mode was entered (jd_debug_b asserted from Nexus1) and EOC = 2'b01
<i>nex_evti_b</i>	I	Event In (<i>nex_evti_b</i>) is an input which, when asserted, initiates one of two events based on the EIC bits in the DC1 Register (if the Nexus2+ module is enabled at reset): - Program Trace and Data Trace synchronization messages (provided Program Trace and Data Trace are enabled and EIC = 2'b00). - Debug request to Zen Nexus1 module (provided EIC = 2'b01 and this feature is implemented).

The Nexus auxiliary port arbitration pins are used when the Nexus 2+ module is implemented in a multi-Nexus SoC which shares a single auxiliary output port. The arbitration is controlled by an SoC level Nexus Port Control module (NPC). Refer to [Section 9.13, "Auxiliary Port Arbitration,"](#) for detail on Nexus port arbitration.

Table 9-25. Nexus Port Arbitration Signals

Nexus Port Arbitration Pins	Input/Output	Description of Arbitration Pins
<i>nex_aux_req</i> [1:0]	O	Nexus Auxiliary Request (<i>nex_aux_req</i> [1:0]) output signals indicate to an SoC level Nexus arbiter a request for access to the shared Nexus auxiliary port in a multi-Nexus implementation. The priority encodings are determined by how many messages are currently in the message queues (see Table 9-23).
<i>nex_aux_busy</i>	O	Nexus Auxiliary Busy (<i>nex_aux_busy</i>) is an output signal to an SoC level Nexus arbiter indicating that the Nexus 2+ module is currently transmitting its message after being granted the Nexus auxiliary port.
<i>npc_aux_grant</i>	I	Nexus Auxiliary Grant (<i>npc_aux_grant</i>) is an input from the SoC level Nexus Port Controller (NPC) that the auxiliary port has been granted to the Nexus 2+ module to transmit its message.
<i>ext_multi_nex_sel</i>	I	Multi-Nexus Select (<i>ext_multi_nex_sel</i>) is a static signal indicating that the Nexus 2+ module is implemented within a multi-Nexus environment. If set, port control and arbitration is controlled by the SoC level arbitration module (NPC).

9.11.2 Pin Protocol

The protocol for the Zen processor transmitting messages via the auxiliary pins is accomplished with the MSEO pin function outlined in [Table 9-26](#). Both single and dual pin cases are shown.

nex_mseo_b[1:0] is used to signal the end of variable-length packets, and not fixed length packets. *nex_mseo_b*[1:0] is sampled on the rising edge of the Nexus 2+ clock (*nex_mcko*).

Table 9-26. MSEO Pin(s) Protocol

nex_mseo_b Function	Single <i>nex_mseo_b</i> data (serial)	Dual <i>nex_mseo_b</i>[1:0] data
Start of message	1-1-0	11-00
End of message	0-1-1-(more 1's)	00 (or 01)-11-(more 1's)
End of variable length packet	0-1-0	00-01
Message transmission	0's	00's
Idle (no message)	1's	11's

Figure 9-31 illustrates the state diagram for single pin MSEO transfers.




Figure 9-31. Single Pin MSEO Transfers

Note that the “End Message” state does not contain valid data on the `nex_mdo[n:0]` pins. Also, It is not possible to have two consecutive “End Packet” messages. This implies the minimum packet size for a variable length packet is 2x the number of `nex_mdo[n:0]` pins. This ensures that a false end of message state is not entered by emitting two consecutive ‘1’s on the `nex_mseo_b` pin before the actual end of message.

Figure 9-32 illustrates the state diagram for dual pin MSEO transfers.




Figure 9-32. Dual Pin MSEO Transfers

The dual pin MSEO option is more robust than the single pin option. Termination of the current message may immediately be followed by the start of the next message on the consecutive clocks. An extra clock to end the message is not necessary as with the one MSEO pin option. The dual pin option also allows for consecutive “End Packet” states. This can be an advantage when small, variable sized packets are transferred.

NOTE

The “End Message” state may also indicate the end of a variable-length packet as well as the end of the message when using the dual pin option.

9.12 Rules for Output Messages

Zen based Class 3 compliant embedded processors must provide messages via the auxiliary port in a consistent manner as described below:

- A variable-sized packet within a message must end on a port boundary.
- A variable-sized packet may start within a port boundary only when following a fixed length packet. (If two variable-sized packets end and start on the same clock, it is impossible to know which bit is from the last packet and which bit is from the next packet.)
- Whenever a variable-length packet is sized such that it does not end on a port boundary, it is necessary to extend and zero fill the remaining bits after the highest-order bit so that it can end on a port boundary.

For example, if the *nex_mdo*[n:0] port is 2 bits wide, and the unique portion of an indirect address TCODE is 5 bits, then the remaining 1 bit of *nex_mdo*[n:0] must be packed with a 0.

9.13 Auxiliary Port Arbitration

In a multi-Nexus environment, the Nexus 2+ module must arbitrate for the shared Nexus port at the SoC level. The request scheme is implemented as a 2-bit request with various levels of priority. The priority levels are defined in Table 9-27 below. The Nexus 2+ module receives a 1-bit grant signal (*npc_aux_grant*) from the SoC level arbiter. When a grant is received, the Nexus 2+ module begins transmitting its message following the protocol outlined in Section 9.11.2. The Nexus 2+ module maintains control of the port by asserting the *nex_aux_busy* signal, until the MSEO state machine reaches the “End Message” state.

Table 9-27. MDO Request Encodings

Request Level	MDO Request Encoding (<i>nex_aux_req</i> [1:0])	Condition of Queue
No Request	00	No message to send
Low Priority	01	Message queue less than 1/2 full
—	10	Reserved
High Priority	11	Message queue 1/2 full or more

9.14 Examples

The following are examples of Program Trace and Data Trace Messages.

Table 9-28 illustrates an example Indirect Branch Message with 2 MDO / 1MSEO configuration.

[Table 9-29](#) illustrates the same example with an 8 MDO / 2 MSEO configuration.

Note that T0 and S0 are the least significant bits where:

- Tx = TCODE number (fixed)
- Sx = Source processor (fixed)
- Ix = Number of instructions (variable)
- Ax = Unique portion of the address (variable)

Note that during clock 12, the *nex_mdo*[n:0] pins are ignored in the single MSEO case.

Table 9-28. Indirect Branch Message Example (2 MDO/1 MSEO)

Clock	<i>nex_mdo</i> [1:0]		<i>nex_mseo_b</i>	State
0	X	X	1	Idle (or end of last message)
1	T1	T0	0	Start Message
2	T3	T2	0	Normal Transfer
3	T5	T4	0	Normal Transfer
4	S1	S0	0	Normal Transfer
5	S3	S2	0	Normal Transfer
6	I1	I0	0	Normal Transfer
7	I3	I2	0	Normal Transfer
8	I5	I4	1	End Packet
9	A1	A0	0	Normal Transfer
10	A3	A2	0	Normal Transfer
11	A5	A4	0	Normal Transfer
12	A7	A6	1	End Packet
13	0	0	1	End Message
14	T1	T0	0	Start Message

Table 9-29. Indirect Branch Message Example (8 MDO/2 MSEO)

Clock	<i>nex_mdo</i> [7:0]								<i>nex_mseo_b</i> [1:0]		State
0	X	X	X	X	X	X	X	X	1	1	Idle (or end of last message)
1	S1	S0	T5	T4	T3	T2	T1	T0	0	0	Start Message
2	I5	I4	I3	I2	I1	I0	S3	S2	0	1	End Packet
3	A7	A6	A5	A4	A3	A2	A1	A0	1	1	End Packet/End Message
4	S1	S0	T5	T4	T3	T2	T1	T0	0	0	Start Message

Table 9-30 and Table 9-31 illustrate examples of Direct Branch Messages: one with 2 MDO / 1 MSEO, and one with 8 MDO / 2 MSEO.

Note that T0 and I0 are the least significant bits where:

- Tx = TCODE number (fixed)
- Sx = Source processor (fixed)
- Ix = Number of Instructions (variable)

Table 9-30. Direct Branch Message Example (2 MDO/1 MSEO)

Clock	nex_mdo[1:0]		nex_mseo_b	State
0	X	X	1	Idle (or end of last message)
1	T1	T0	0	Start Message
2	T3	T2	0	Normal Transfer
3	T5	T4	0	Normal Transfer
4	S1	S0	0	Normal Transfer
5	S3	S2	0	Normal Transfer
6	I1	I0	1	End Packet
7	0	0	1	End Message

Table 9-31. Direct Branch Message Example (8 MDO/2 MSEO)

Clock	nex_mdo[7:0]								nex_mseo_b[1:0]		State
0	X	X	X	X	X	X	X	X	1	1	Idle (or end of last message)
1	S1	S0	T5	T4	T3	T2	T1	T0	0	0	Start Message
2	0	0	0	0	I1	I0	S3	S2	1	1	End Packet/End Message
3	S1	S0	T5	T4	T3	T2	T1	T0	0	0	Start Message

Table 9-32 illustrates an example Data Write Message with 8 MDO / 1 MSEO configuration, and Table 9-33 illustrates the same DWM with 8 MDO / 2 MSEO configuration

Note that T0, A0, D0 are the least significant bits where:

- Tx = TCODE number (fixed)
- Sx = Source processor (fixed)
- Zx = Data size (fixed)
- Ax = Unique portion of the address (variable)
- Dx = Write data (variable—8, 16 or 32-bit)

Table 9-32. Data Write Message Example (8 MDO/1 MSEO)

Clock	nex_mdo[7:0]								nex_mseo_b	State
0	X	X	X	X	X	X	X	X	1	Idle (or end of last message)
1	S1	S0	T5	T4	T3	T2	T1	T0	0	Start Message
2	A2	A1	A0	Z2	Z1	Z0	S3	S2	1	End Packet
3	D7	D6	D5	D4	D3	D2	D1	D0	0	Normal Transfer
4	0	0	0	0	0	0	0	0	1	End Packet
5	0	0	0	0	0	0	0	0	1	End Message

Table 9-33. Data Write Message Example (8 MDO/2 MSEO)

Clock	<i>nex_mdo</i> [7:0]								<i>nex_mseo_b</i> [1:0]		State
0	X	X	X	X	X	X	X	X	1	1	Idle (or end of last message)
1	S1	S0	T5	T4	T3	T2	T1	T0	0	0	Start Message
2	A2	A1	A0	Z2	Z1	Z0	S3	S2	0	1	End Packet
3	D7	D6	D5	D4	D3	D2	D1	D0	1	1	End Packet/ End Message

9.15 IEEE 1149.1 (JTAG) RD/WR Sequences

This section contains example JTAG/OnCE sequences used to access resources.

9.15.1 JTAG Sequence for Accessing Internal Nexus Registers

Table 9-34. Accessing Internal Nexus 2+ Registers via JTAG/OnCE

Step #	TMS Pin	Description
1	1	IDLE -> SELECT-DR_SCAN
2	0	SELECT-DR_SCAN -> CAPTURE-DR (Nexus Command Register value loaded in shifter)
3	0	CAPTURE-DR -> SHIFT-DR
4	0	(7) TCK clocks issued to shift in direction (rd/wr) bit and first 6 bits of Nexus reg. addr.
5	1	SHIFT-DR -> EXIT1-DR (7th bit of Nexus reg. shifted in)
6	1	EXIT1-DR -> UPDATE-DR (Nexus shifter is transferred to Nexus Command Register)
7	1	UPDATE-DR -> SELECT-DR_SCAN
8	0	SELECT-DR_SCAN -> CAPTURE-DR (Register value is transferred to Nexus shifter)
9	0	CAPTURE-DR -> SHIFT-DR
10	0	(31) TCK clocks issued to transfer register value to TDO pin while shifting in TDI value
11	1	SHIFT-DR -> EXIT1-DR (MSB of value is shifted in/out of shifter)
12	1	EXIT1-DR -> UPDATE -DR (if access is write, shifter is transferred to register)
13	0	UPDATE-DR -> RUN-TEST/IDLE (transfer complete—Nexus controller to Reg. Select state)

9.15.2 JTAG Sequence for Read Access of Memory-Mapped Resources

Table 9-35. Accessing Memory-Mapped Resources (Reads)

Step #	TCLK Clocks	Description
1	13	Nexus Command = write to Read/Write Access Address Register (RWA)
2	37	Write RWA (initialize starting read address—data input on TDI)
3	13	Nexus Command = write to Read/Write Control/Status Register (RWCS)
4	37	Write RWCS (initialize read access mode and CNT value—data input on TDI)
5	—	Wait for falling edge of <i>nex_rdy_b</i> pin
6	13	Nexus Command = read Read/Write Access Data Register (RWD)
7	37	Read RWD (data output on TDO)
8	—	If CNT > 0, go back to Step #5

9.15.3 JTAG Sequence for Write Access of Memory-Mapped Resources

Table 9-36. Accessing Memory-Mapped Resources (Writes)

Step #	TCLK Clocks	Description
1	13	Nexus Command = write to Read/Write Access Control/Status Register (RWCS)
2	37	Write RWCS (initialize write access mode and CNT value—data input on TDI)
3	13	Nexus Command = write to Read/Write Address Register (RWA)
4	37	Write RWA (initialize starting write address—data input on TDI)
5	13	Nexus Command = read Read/Write Access Data Register (RWD)
6	37	Write RWD (data output on TDO)
7	—	Wait for falling edge of <i>nex_rdy_b</i> pin
8	—	If CNT > 0, go back to Step #5



Appendix A

Register Summary





Figure A-1. e200 User Mode Registers



BTB Register

BTB Control¹


BUCSR

SPR 1013

¹ - These e200-specific registers may not be supported by other PowerPC processors

² - Optional registers defined by the PowerPC Book-E architecture

Figure A-2. e200z0h Supervisor Mode Registers



¹ - These e200-specific registers may not be supported by other PowerPC processors

² - Optional registers defined by the PowerPC Book-E architecture

Figure A-3. e200z0 Supervisor Mode Registers

e200z0h	0	UCLE	Allocated	0	WE	CE	0	EE	PR	FP	ME	FE0	0	DE	FE1	0	IS	DS	0	RI	0
e200z0	0	UCLE	Allocated	0	WE	CE	0	EE	PR	FP	ME	FE0	0	DE	FE1	0	IS	DS	0		
																					0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Figure A-4. Machine State Register (MSR)

0	ID
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31	

Figure A-5. Processor ID Register (PIR)

1	0	0	0	0	0	0	1	0	0	1	0	Version				MBG Reserved								Major Rev				MBG ID			
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-6. Processor Version Register (PVR)

System Version																															
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-7. System Version Register (SVR)

SO	OV	CA	0																						Bytecnt						
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-8. Integer Exception Register (XER)

0				PIL	PPR	PTR	FP	ST	0	DLK	ILK	AP	PUO	BO	PIE	0							EFP	0	VLEMI	0			MIF	XTE	
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-9. Exception Syndrome Register (ESR)

e200z0h	MCP	0	CP_PERR	CPERR	EXCP_ERR	0	NMI	0	BUS_IRERR	BUS_DRERR	BUS_WRERR	0
e200z0	MCP	0	CP_PERR	CPERR	EXCP_ERR	0	0	0	BUS_IRERR	BUS_DRERR	BUS_WRERR	0

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Figure A-10. Machine Check Syndrome Register (MCSR)

EMCP	0	BPRED	DOZE	NAP	SLEEP	0	ICR	NHR	0	TBEN	SEL_TBCK	DCLREE	DCLRCE	CICLRDE	MCCLRDE	DAPUEN	0
------	---	-------	------	-----	-------	---	-----	-----	---	------	----------	--------	--------	---------	---------	--------	---

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Figure A-11. Hardware Implementation-Dependent Register 0 (HID0)

0	SYSCTL	ATS	0
---	--------	-----	---

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Figure A-12. Hardware Implementation-Dependent Register 1 (HID1)

0	BBFI	0	BPEN
---	------	---	------

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Figure A-13. Branch Unit Control and Status Register (BUCSR)

EDM	IDM	RST	ICMP	BRT	IRPT	TRAP	IAC1	IAC2	IAC3	IAC4	DAC1	DAC2	RET	0	DEVT1	DEVT2	0	CIRPT	CRET	0
-----	-----	-----	------	-----	------	------	------	------	------	------	------	------	-----	---	-------	-------	---	-------	------	---

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Figure A-14. DBCR0 Register

IAC1US	IAC1ER	IAC2US	IAC2ER	IAC12M	0	IAC3US	IAC3ER	IAC4US	IAC4ER	IAC34M	0
--------	--------	--------	--------	--------	---	--------	--------	--------	--------	--------	---

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Figure A-15. DBCR1 Register

Register Summary

DAC1US	DAC1ER	DAC2US	DAC2ER	DAC12M	DAC1LNK	DAC2LNK	0																								
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-16. DBCR2 Register

IDE	UDE	MRR		ICMP	BRT	IRPT	TRAP	IAC1	IAC2	IAC3	IAC4	DAC1R	DAC1W	DAC2R	DAC2W	RET	0				DEVT1	DEVT2	0			CIRPT	CRET	VLES	0			
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	

Figure A-17. DBSR Register

MCLK	ERR	CHKSTOP	RESET	HALT	STOP	DEBUG	WAIT	0	1
0	1	2	3	4	5	6	7	8	9

Figure A-18. OnCE Status Register

R/W	GO	EX	RS[0:6]						
0	1	2	3	4	5	6	7	8	9

Figure A-19. OnCE Command Register

0																DMDIS	0	DW	DI	DM	DG	DE	0					WKUP	FDB	DR	
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-20. OnCE Control Register




Figure A-21. CPU Scan Chain Register (CPUSCR)

*															WAITING	PCOFST		PCINV	FFRA	IRSTAT0	IRSTAT1	IRSTAT2	IRSTAT3	IRSTAT4	IRSTAT5	IRSTAT6	IRSTAT7	IRSTAT8	IRSTAT9		
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-22. Control State Register (CTL)

CARCH		CWPA	CFAHA	CFIWA	0		CBSIZE	CREPL		CLA	CPA	CNWAY								CSIZE											
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-23. L1 Cache Configuration Register 0 (L1CFG0)

Register Summary

0																	NPIDS			PIDSIZE				0		NTLBS		MAVN			
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Figure A-24. MMU Configuration Register (MMUCFG)

Glossary

The glossary contains an alphabetical list of terms, phrases, and abbreviations used in this reference manual.

A

Architecture. A detailed specification of requirements for a processor or computer system. It does not specify details of how the processor or computer system must be implemented; instead it provides a template for a family of compatible *implementations*.

Atomic access. A bus access that attempts to be part of a read-write operation to the same address uninterrupted by any other access to that address (the term refers to the fact that the transactions are indivisible). The Power Architecture technology implements atomic accesses through the **lwarx/stwcx** instruction pair.

Autobaud. The process of determining a serial data rate by timing the width of a single bit.

B

Beat. A single state on the bus interface that may extend across multiple bus cycles. A transaction can be composed of multiple address or data *beats*.

Big-Endian. A byte-ordering method in memory where the address *n* of a word corresponds to the *most significant byte*. In an addressed memory word, the bytes are ordered (left to right) 0, 1, 2, 3, with 0 being the *most significant byte*. See *Little-Endian*.

Boundedly undefined. A characteristic of certain operation results that are not rigidly prescribed by the Power Architecture technology. Boundedly-undefined results for a given operation may vary among implementations and between execution attempts in the same implementation.

Although the architecture does not prescribe the exact behavior for when results are allowed to be boundedly undefined, the results of executing instructions in contexts where results are allowed to be boundedly undefined are constrained to ones that could have been achieved by executing an arbitrary sequence of defined instructions, in valid form, starting in the state the machine was in before attempting to execute the given instruction.

Breakpoint. A programmable event that forces the core to take a breakpoint exception.

Burst. A multiple-beat data transfer whose total size is typically equal to a cache block.

Bus clock. Clock that causes the bus state transitions.

Bus master. The owner of the address or data bus; the device that initiates or requests the transaction.

C

Cache. High-speed memory containing recently accessed data or instructions (subset of main memory).

Cache block. A small region of contiguous memory that is copied from memory into a [cache](#). The size of a cache block may vary among processors; the maximum block size is one [page](#). In Power Architecture processors, [cache coherency](#) is maintained on a cache-block basis. Note that the term ‘cache block’ is often used interchangeably with ‘cache line.’

Cache coherency. An attribute wherein an accurate and common view of memory is provided to all devices that share the same memory system. Caches are coherent if a processor performing a read from its cache is supplied with data corresponding to the most recent value written to memory or to another processor’s cache.

Cache flush. An operation that removes from a cache any data from a specified address range. This operation ensures that any modified data within the specified address range is written back to main memory. This operation is generated typically by a Data Cache Block Flush (**dcbf**) instruction.

Caching-inhibited. A memory update policy in which the [cache](#) is bypassed and the load or store is performed to or from main memory.

Cast out. A [cache block](#) that must be written to memory when a cache miss causes a cache block to be replaced.

Changed bit. One of two *page history bits* found in each [page table entry](#) (PTE). The processor sets the changed bit if any store is performed into the *page*. See also [Page access history bits](#) and [Referenced bit](#).

Clean. An operation that causes a cache block to be written to memory, if modified, and then left in a valid, unmodified state in the cache.

Clear. To cause a bit or bit field to register a value of zero. See also [Set](#).

Completer. In PCI-X, a completer is the device addressed by a transaction (other than a split completion transaction). If a target terminates a transaction with a split response, the completer becomes the initiator of the subsequent split completion.

Context synchronization. An operation that ensures that all instructions in execution complete past the point where they can produce an *exception*, that all instructions in execution complete in the context in which they began execution, and that all subsequent instructions are [fetched](#) and executed in the new context. Context synchronization may result from executing specific instructions (such as **isync** or **rfi**) or when certain events occur (such as an exception).

Copy-back operation. A cache operation in which a cache line is copied back to memory to enforce cache coherency. Copy-back operations consist of snoop push-out operations and cache cast-out operations.

D **Direct-mapped cache.** A cache in which each main memory address can appear in only one location within the cache; operates more quickly when the memory request is a cache hit.

Double data rate. Memory that allows data transfers at the start and end of a clock cycle, thereby doubling the data rate.

E **Effective address (EA).** The 32-bit address specified for a load, store, or an instruction fetch. This address is then submitted to the MMU for translation to either a *physical memory* or an I/O address.

Exclusive state. MEI state (E) in which only one caching device contains data that is also in system memory.

F **Fetch.** Retrieving instructions from either the cache or main memory and placing them into the instruction queue.

Flush. An operation that causes a cache block to be invalidated and the data, if modified, to be written to memory.

Frame-check sequence (FCS). Specifies the standard 32-bit cyclic redundancy check (CRC) obtained using the standard CCITT-CRC polynomial on all fields except the preamble, SFD, and CRC.

G **General-purpose register (GPR).** Any of the 32 registers in the general-purpose register file. These registers provide the source operands and destination results for all integer data manipulation instructions. Integer load instructions move data from memory to GPRs and store instructions move data from GPRs to memory.

Gigabit media-independent interface (GMII) sublayer. Sublayer that provides a standard interface between the MAC layer and the physical layer for 1000-Mbps operation. It isolates the MAC layer and the physical layer, enabling the MAC layer to be used with various implementations of the physical layer.

Guarded. The guarded attribute pertains to out-of-order execution. When a page is designated as guarded, instructions and data cannot be accessed out-of-order.

H **Harvard architecture.** An architectural model featuring separate caches and other memory management resources for instructions and data.

I

Illegal instructions. A class of instructions that are not implemented for a particular processor. These include instructions not defined by the architecture. In addition, for 32-bit implementations, instructions that are defined only for 64-bit implementations are considered to be illegal instructions. For 64-bit implementations, instructions that are defined only for 32-bit implementations are considered to be illegal instructions.

Implementation. A particular processor that conforms to the architecture, but may differ from other architecture-compliant implementations (for example, in design, feature set, and implementation of *optional* features).

Imprecise exception. A type of *synchronous exception* that is allowed not to adhere to the precise exception model (see *Precise exception*). The Power Architecture technology allows only floating-point exceptions to be handled imprecisely.

Inbound ATMU windows. Mappings that perform address translation from the external address space to the local address space, attach attributes and transaction types to the transaction, and map the transaction to its target interface.

In-order. An aspect of an operation that adheres to a sequential model. An operation is said to be performed *in-order* if, at the time that it is performed, it is known to be required by the sequential execution model.

Integer unit. An execution unit in the core responsible for executing integer instructions.

Instruction latency. The total number of clock cycles necessary to execute an instruction and make ready the results of that instruction.

K

Kill. An operation that causes a *cache block* to be invalidated without writing any modified data to memory.

L

L2 cache. Level-2 cache. See *Secondary cache*.

Latency. The number of clock cycles necessary to execute an instruction and make ready the results of that execution for a subsequent instruction.

Least significant bit (lsb). The bit of least value in an address, register, field, data element, or instruction encoding.

Least significant byte (LSB). The byte of least value in an address, register, data element, or instruction encoding.

Little-Endian. A byte-ordering method in memory where the address *n* of a word corresponds to the *least significant byte*. In an addressed memory word, the bytes are ordered (left to right) 3, 2, 1, 0, with 3 being the *most significant byte*. See *Big-Endian*.

Local access window. Mapping used to translate a region of memory to a particular target interface, such as the DDR SDRAM controller or the PCI controller. The local memory map is defined by a set of eight local access windows. The size of each window can be configured from 4 Kbytes to 2 Gbytes.

M

Media access control (MAC) sublayer. Sublayer that provides a logical connection between the MAC and its peer station. Its primary responsibility is to initialize, control, and manage the connection with the peer station.

Medium-dependent interface (MDI) sublayer. Sublayer that defines different connector types for different physical media and PMD devices.

Media-independent interface (MII) sublayer. Sublayer that provides a standard interface between the MAC layer and the physical layer for 10/100-Mbps operations. It isolates the MAC layer and the physical layer, enabling the MAC layer to be used with various implementations of the physical layer.

Memory access ordering. The specific order in which the processor performs load and store memory accesses and the order in which those accesses complete.

Memory coherency. An aspect of caching in which it is ensured that an accurate view of memory is provided to all devices that share system memory.

Memory consistency. Refers to agreement of levels of memory with respect to a single processor and system memory (for example, on-chip cache, secondary cache, and system memory).

Memory management unit (MMU). The functional unit that is capable of translating an *effective (logical) address* to a physical address, providing protection mechanisms, and defining caching methods.

Memory-mapped accesses. Accesses whose addresses use the page or block address translation mechanisms provided by the MMU and that occur externally with the bus protocol defined for memory.

Modified/exclusive/invalid (MEI). *Cache coherency* protocol used to manage caches on different devices that share a memory system. Note that the Power Architecture technology does not specify the implementation of an MEI protocol to ensure cache coherency.

Modified state. MEI state (M) in which one, and only one, caching device has the valid data for that address. The data at this address in external memory is not valid.

Most significant bit (msb). The highest-order bit in an address, registers, data element, or instruction encoding.

Most significant byte (MSB). The highest-order byte in an address, registers, data element, or instruction encoding.

N

NaN. An abbreviation for not a number; a symbolic entity encoded in floating-point format. There are two types of NaNs—signaling NaNs and quiet NaNs.

No-op. No-operation. A single-cycle operation that does not affect registers or generate bus activity.

O

OCeaN (on-chip network). Non-blocking crossbar switch fabric. Enables full duplex port connections at 128 Gb/s concurrent throughput and independent per port transaction queuing and flow control. Permits high bandwidth, high performance, as well as the execution of multiple data transactions.

Outbound ATMU windows. Mappings that perform address translations from local 32-bit address space to the address spaces of RapidIO or PCI/PCI-X RapidIO, which may be much larger than the local space. Outbound ATMU windows also map attributes such as transaction type or priority level.

P

Packet. A unit of binary data that can be routed through a network. Sometimes packet is used to refer to the frame plus the preamble and start frame delimiter (SFD).

Page. A region in memory. The OEA defines a page as a 4-Kbyte area of memory aligned on a 4-Kbyte boundary.

Page access history bits. The [changed](#) and [referenced](#) bits in the PTE keep track of the access history within the page. The referenced bit is set by the MMU whenever the page is accessed for a read or write operation. The changed bit is set when the page is stored into. See [Changed bit](#) and [Referenced bit](#).

Page fault. A page fault is a condition that occurs when the processor attempts to access a memory location that does not reside within a *page* not currently resident in [physical memory](#). A page fault exception condition occurs when a matching, valid [page table entry](#) (PTE[V] = 1) cannot be located.

Page table. A table in memory is comprised of [page table entries](#), or PTEs. It is further organized into eight PTEs per PTEG (page table entry group). The number of PTEGs in the page table depends on the size of the page table (as specified in the SDR1 register).

Page table entry (PTE). Data structures containing information used to translate [effective address](#) to physical address on a 4-Kbyte page basis. A PTE consists of 8 bytes of information in a 32-bit processor and 16 bytes of information in a 64-bit processor.

Physical coding sublayer (PCS). Sublayer responsible for encoding and decoding data stream to and from the MAC sublayer. Medium (1000BASEX) 8B/10B coding is used for fiber. Medium (1000BASET) 8B1Q coding is used for unshielded twisted pair (UTP).

Physical medium attachment (PMA) sublayer. Sublayer responsible for serializing code groups into a bit stream suitable for serial bit-oriented physical devices (SerDes) and vice versa. Synchronization is also performed for proper data decoding in this sublayer. The PMA sits between the PCS and the PMD sublayers. For fiber medium (1000BASEX) the interface on the PMD side of the PMA is a one-bit 1250-MHz signal, while on the PMA PCS side, the interface is a ten-bit interface (TBI) at 125 MHz. The TBI is an alternative to the GMII interface. If the TBI is used, the gigabit Ethernet controller must be capable of performing the PCS function. For UTP medium, the PMD interface side of the PMA consists of four pair of 62.5-MHz PAM5 encoded signals, while the PCS side provides the 1250-Mbps input to an 8B1Q4 PCS.

Physical medium dependent (PMD) sublayer. Sublayer responsible for signal transmission. The typical PMD functionality includes amplifier, modulation, and wave shaping. Different PMD devices may support different media.

Physical memory. The actual memory that can be accessed through the system's memory bus.

Pipelining. A technique that breaks operations, such as instruction processing or bus transactions, into smaller distinct stages or tenures (respectively) so that a subsequent operation can begin before the previous one has completed.

Precise exceptions. A category of exception for which the pipeline can be stopped so instructions that preceded the faulting instruction can complete and subsequent instructions can be flushed and redispached after exception handling has completed. See [Imprecise exceptions](#).

Primary opcode. The most-significant 6 bits (bits 0–5) of the instruction encoding that identifies the type of instruction.

Program order. The order of instructions in an executing program. More specifically, this term is used to refer to the original order in which program instructions are fetched into the instruction queue from the cache.

Protection boundary. A boundary between [protection domains](#).

Protection domain. A protection domain is a segment, a virtual page, a BAT area, or a range of unmapped effective addresses. It is defined only when the appropriate relocate bit in the MSR (IR or DR) is 1.

Q

Quad word. A group of 16 contiguous locations starting at an address divisible by 16.

Quiesce. To come to rest. The processor is said to quiesce when an exception is taken or a **sync** instruction is executed. The instruction stream is stopped at the decode stage and executing instructions are allowed to complete to create a controlled context for instructions that may be affected by out-of-order, parallel execution. See [Context synchronization](#).

R

rA. The **rA** instruction field is used to specify a GPR to be used as a source or destination.

rB. The **rB** instruction field is used to specify a GPR to be used as a source.

rD. The **rD** instruction field is used to specify a GPR to be used as a destination.

rS. The **rS** instruction field is used to specify a GPR to be used as a source.

RapidIO. High-performance, packet-switched, interconnect architecture that provides reliability, increased bandwidth, and faster bus speeds in an intra-system interconnect. Designed to be compatible with integrated communications processors, host processors, and networking digital signal processors,

Reconciliation sublayer. Sublayer that maps the terminology and commands used in the MAC layer into electrical formats appropriate for the physical layer entities.

Record bit. Bit 31 (or the Rc bit) in the instruction encoding. When it is set, updates the condition register (CR) to reflect the result of the operation.

Reduced instruction set computing (RISC). An *architecture* characterized by fixed-length instructions with nonoverlapping functionality and by a separate set of load and store instructions that perform memory accesses.

Referenced bit. One of two *page history bits* found in each *page table entry*. The processor sets the *referenced bit* whenever the page is accessed for a read or write. See also *Page access history bits*.

Requester. In PCI-X, a requester is an initiator that first introduces a transaction into the PCI-X domain. If a transaction is terminated with a split response, the requester becomes the target of the subsequent split completion.

Reservation. The processor establishes a reservation on a *cache block* of memory space when it executes an **lwarx** instruction to read a memory semaphore into a GPR.

Reservation station. A buffer between the dispatch and execute stages that allows instructions to be dispatched even though the results of instructions on which the dispatched instruction may depend are not available.

S

Secondary cache. A cache memory that is typically larger and has a longer access time than the primary cache. A secondary cache may be shared by multiple devices. Also referred to as L2, or level-2, cache.

Sequence. In PCI-X, a sequence is one or more transactions associated with carrying out a single logical transfer by a requester. Each transaction in the same sequence carries the same unique sequence ID.

Set (v). To write a nonzero value to a bit or bit field; the opposite of *clear*. The term ‘set’ may also be used to generally describe the updating of a bit or bit field.

Set (n). A subdivision of a [cache](#). Cacheable data can be stored in a given location in one of the sets, typically corresponding to its lower-order address bits. Because several memory locations can map to the same location, cached data is typically placed in the set whose [cache block](#) corresponding to that address was used least recently. See [Set associative](#).

Set associative. Aspect of cache organization in which the cache space is divided into sections, called [sets](#). The cache controller associates a particular main memory address with the contents of a particular set, or region, within the cache.

Slave. The device addressed by a master device. The slave is identified in the address tenure and is responsible for supplying or latching the requested data for the master during the data tenure.

Snooping. Monitoring addresses driven by a bus master to detect the need for coherency actions.

Snoop push. Response to a snooped transaction that hits a modified cache block. The cache block is written to memory and made available to the snooping device.

Stall. An occurrence when an instruction cannot proceed to the next stage.

Sticky bit. A bit that when [set](#) must be cleared explicitly.

Superscalar machine. A machine that can issue multiple instructions concurrently from a conventional linear instruction stream.

Supervisor mode. The privileged operation state of a processor. In supervisor mode, software, typically the operating system, can access all control registers and can access the supervisor memory space, among other privileged operations.

Synchronization. A process to ensure that operations occur strictly [in order](#). See [Context synchronization](#).

Synchronous exception. An *exception* that is generated by the execution of a particular instruction or instruction sequence. There are two types of synchronous exceptions, [precise](#) and [imprecise](#).

System memory. The physical memory available to a processor.

T

Tenure. The period of bus mastership. There can be separate address bus tenures and data bus tenures.

Throughput. The measure of the number of instructions that are processed per clock cycle.

Time-division multiplex (TDM). A single serial channel used by several channels taking turns.

Transaction. A complete exchange between two bus devices. A transaction is typically comprised of an address tenure and one or more data tenures, which may overlap or occur separately from the address tenure. A transaction may be minimally comprised of an address tenure only.

Transfer termination. Signal that refers to both signals that acknowledge the transfer of individual beats (of both single-beat transfer and individual beats of a burst transfer) and to signals that mark the end of the tenure.

Translation lookaside buffer (TLB). A cache that holds recently-used *page table entries*.

U

User mode. The operating state of a processor used typically by application software. In user mode, software can access only certain control registers and can access only user memory space. No privileged operations can be performed. Also referred to as problem state.

V

Virtual address. An intermediate address used in the translation of an *effective address* to a physical address.

Virtual memory. The address space created using the memory management facilities of the processor. Program access to *virtual memory* is possible only when it coincides with *physical memory*.

W

Way. A location in the cache that holds a cache block, its tags, and status bits.

Word. A 32-bit data element.

Write-back. A cache memory update policy in which processor write cycles are directly written only to the cache. External memory is updated only indirectly, for example, when a modified cache block is *cast out* to make room for newer data.

Write-through. A cache memory update policy in which all processor write cycles are written to both the cache and memory.

Index

A

Alignment
exception, 5-15

B

Branch prediction
instruction model, 3-3
Breakpoint request signal, 8-22
Bus
transfers, 6-14

C

Control register (CTL), 8-34
CR (condition register)
CR, description, 2-4
CTR register, 2-5
Cycle
Bus transfer, 6-14

D

DAR (data address register), 2-5
Debug
Breakpoint request, 8-22
Mode, 8-20, 8-31
Registers, 8-34, 8-37, 8-38
Reset interaction, 8-31
Scan chain, 8-31, 8-32, 8-34, 8-37, 8-38
Signals, 8-20, 8-21, 8-22
Debug mode signal, 8-22
Debug request signal, 8-20
DSI exception, 5-13, 5-17

E

Exceptions
alignment exception, 5-15
DSI exception, 5-13, 5-17
enabling and disabling exceptions, 5-26
machine check exception, 5-11
program exception, 5-15
register settings
MSR, 2-8, 2-12, 5-4, 5-5
reset exception, 5-20
returning from an exception handler, 5-27
summary table, 5-2
system call exception, 5-17
terminology, 5-1

G

GPRn (general-purpose registers), 2-4

I

Instructions
isync, 5-28
listed by mnemonic, 3-4–??
rfi, 5-27
stwcx., 5-28
isync, 5-28

L

LR (link register), 2-5

M

Machine check exception, 5-11
MSR (machine state register)
bit settings, 2-8, 2-12, 5-4, 5-5

P

Process switching, 5-28
Processor status register (PSR), 8-38
Program counter, 8-37
Program exception, 5-15
PVR (processor version register), 2-5

R

Registers
reset settings, 2-23
supervisor-level
DAR, 2-5
MSR, 2-5
PVR, 2-5
SRR0/SRR1, 2-6, 2-7
user-level
CR, 2-4
CTR, 2-5
GPRn, 2-4
LR, 2-5
XER, 2-4
Reset
Debug interaction, 8-31
reset exception, 5-20
rfi, 5-27

S

Scan chain, debug, 8-32

Signals

- Breakpoint, 8-22

- Debug, 8-20, 8-21

- Debug acknowledge, 8-20

- Debug mode, 8-22

- Debug request, 8-20

- Status, 8-20, 8-21

- Transfer attribute, 6-14

SRR0/SRR1 (status save/restore registers)

- description, 2-6, 2-7

Status signals, 8-20, 8-21

stwcx., 5-28

Synchronization

- execution of rfi, 5-27

System call exception, 5-17

T

Transfer attribute signals, 6-14

W

Write-back bus register (WBRR), 8-37

X

XER register, 2-4



e200z0 and e200z0h Overview	1
Register Model	2
Instruction Model	3
Instruction Pipeline and Execution Timing	4
Interrupts and Exceptions	5
Core Complex Interfaces	6
Power Management	7
Debug Support	8
Nexus 2+ Module	9
Register Summary	A
Glossary	GLO
Index	IND



1	e200z0 and e200z0h Overview
2	Register Model
3	Instruction Model
4	Instruction Pipeline and Execution Timing
5	Interrupts and Exceptions
6	Core Complex Interfaces
7	Power Management
8	Debug Support
9	Nexus 2+ Module
A	Register Summary
GLO	Glossary
IND	Index