

# MIPS32<sup>TM</sup> Architecture For Programmers Volume III: The MIPS32<sup>TM</sup> Privileged Resource Architecture

Document Number: MD00090 Revision 2.00 June 9, 2003

MIPS Technologies, Inc. 1225 Charleston Road Mountain View, CA 94043-1353

Copyright © 2001-2003 MIPS Technologies Inc. All rights reserved.

Copyright © 2001-2003 MIPS Technologies, Inc. All rights reserved.

Unpublished rights (if any) are reserved under the Copyright Laws of the United States of America.

If this document is provided in source format (i.e., in a modifiable form such as in FrameMaker or Microsoft Word format), then its use and distribution is subject to a written agreement with MIPS Technologies, Inc. ("MIPS Technologies"). UNDER NO CIRCUMSTANCES MAY A DOCUMENT PROVIDED IN SOURCE FORMAT BE DISTRIBUTED TO A THIRD PARTY WITHOUT THE EXPRESS WRITTEN CONSENT OF MIPS TECHNOLOGIES.

This document contains information that is proprietary to MIPS Technologies. Any copying, reproducing, modifying, or use of this information (in whole or in part) which is not expressly permitted in writing by MIPS Technologies or a contractually-authorized third party is strictly prohibited. At a minimum, this information is protected under unfair competition and copyright laws. Violations thereof may result in criminal penalties and fines.

MIPS Technologies or any contractually-authorized third party reserves the right to change the information contained in this document to improve function, design or otherwise. MIPS Technologies does not assume any liability arising out of the application or use of this information, or of any error of omission in such information. Any warranties, whether express, statutory, implied or otherwise, including but not limited to the implied warranties of merchantability or fitness for a particular purpose, are excluded. Any license under patent rights or any other intellectual property rights owned by MIPS Technologies or third parties shall be conveyed by MIPS Technologies or any contractually-authorized third party in a separate license agreement between the parties.

The information contained in this document shall not be exported or transferred for the purpose of reexporting in violation of any U.S. or non-U.S. regulation, treaty, Executive Order, law, statute, amendment or supplement thereto.

The information contained in this document constitutes one or more of the following: commercial computer software, commercial computer software documentation or other commercial items. If the user of this information, or any related documentation of any kind, including related technical data or manuals, is an agency, department, or other entity of the United States government ("Government"), the use, duplication, reproduction, release, modification, disclosure, or transfer of this information, or any related documentation of any kind, is restricted in accordance with Federal Acquisition Regulation 12.212 for civilian agencies and Defense Federal Acquisition Regulation Supplement 227.7202 for military agencies. The use of this information by the Government is further restricted in accordance with the terms of the license agreement(s) and/or applicable contract terms and conditions covering this information from MIPS Technologies or any contractually-authorized third party.

MIPS<sup>®</sup>, R3000<sup>®</sup>, R4000<sup>®</sup>, R5000<sup>®</sup> and R10000<sup>®</sup> are among the registered trademarks of MIPS Technologies, Inc. in the United States and certain other countries, and MIPS16<sup>TM</sup>, MIPS16e<sup>TM</sup>, MIPS32<sup>TM</sup>, MIPS64<sup>TM</sup>, MIPS-3D<sup>TM</sup>, MIPS-based<sup>TM</sup>, MIPS II<sup>TM</sup>, MIPS III<sup>TM</sup>, MIPS IV<sup>TM</sup>, MIPS V<sup>TM</sup>, MDMX<sup>TM</sup>, MIPSsim<sup>TM</sup>, MIPSsimCA<sup>TM</sup>, MIPSsimIA<sup>TM</sup>, QuickMIPS<sup>TM</sup>, SmartMIPS<sup>TM</sup>, MIPS Technologies logo, 4K<sup>TM</sup>, 4Kc<sup>TM</sup>, 4Km<sup>TM</sup>, 4Ke<sup>TM</sup>, 4KEc<sup>TM</sup>, 4KEc<sup>TM</sup>, 4KEm<sup>TM</sup>, 4KEp<sup>TM</sup>, 4KS<sup>TM</sup>, 4KS<sup>TM</sup>, 5K<sup>TM</sup>, 5Kc<sup>TM</sup>, 5Kf<sup>TM</sup>, 20K<sup>TM</sup>, 20Kc<sup>TM</sup>, 25Kf<sup>TM</sup>, R4300<sup>TM</sup>, ASMACRO<sup>TM</sup>, ATLAS<sup>TM</sup>, BusBridge<sup>TM</sup>, CoreFPGA<sup>TM</sup>, CoreLV<sup>TM</sup>, EC<sup>TM</sup>, JALGO<sup>TM</sup>, MALTA<sup>TM</sup>, MGB<sup>TM</sup>, PDtrace<sup>TM</sup>, SEAD-2<sup>TM</sup>, SOC-it<sup>TM</sup>, The Pipeline<sup>TM</sup>, and YAMON<sup>TM</sup> are among the trademarks of MIPS Technologies, Inc.

All other trademarks referred to herein are the property of their respective owners.

Template: B1.06, Build with Conditional Tags: 2B ARCH MIPS32

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

Copyright © 2001-2003 MIPS Technologies Inc. All rights reserved.

# Table of Contents

Chapter 1 About This Book	1
1.1 Typographical Conventions	1
1.1.1 Italic Text	1
1.1.2 Bold Text	1
1.1.3 Courier Text	1
1.2 UNPREDICTABLE and UNDEFINED	2
1.2.1 UNPREDICTABLE	2
1.2.2 UNDEFINED	2
1.3 Special Symbols in Pseudocode Notation	2
1.4 For More Information	4
Chapter 2 The MIPS32 Privileged Resource Architecture	
2.1 Introduction	
2.2 The MIPS Coprocessor Model	
2.2.1 CP0 - The System Coprocessor	
2.2.2 CP0 Registers	
Charter 2 MIDE22 Operating Modes	0
2 1 Dahya Mada	
5.1 Debug Mode	
3.2 Kernel Mode	
3.5 Supervisor Mode	
5.4 User Mode	10
2.5.1 64 hit Electing Doint Operations English	10
3.5.1 04-bit FIDE Enable	10
5.5.2 04-bit FFK Eliable	10
Chapter 4 Virtual Memory	11
4.1 Support in Release 1 and Release 2 of the Architecture	11
4.1.1 Virtual Memory	11
4.2 Terminology	11
4.2.1 Address Space	11
4.2.2 Segment and Segment Size	11
4.2.3 Physical Address Size (PABITS)	11
4.3 Virtual Address Spaces	12
4.4 Compliance	
4.5 Access Control as a Function of Address and Operating Mode	
4.6 Address Translation and Cache Coherency Attributes for the kseg0 and kseg1 Segments	
4.7 Address Translation for the kuseg Segment when StatusERL = 1	
4.8 Special Behavior for the kseg3 Segment when DebugDM = 1	
4.9 TLB-Based Virtual Address Translation	
4.9.1 Address Space Identifiers (ASID)	
4.9.2 TLB Organization	17
4.9.3 TLB Initialization	17
4.9.4 Address Translation	19
Chapter 5 Interrupts and Exceptions	
5.1 Interrupts	23
5.1.1 Interrupt Modes	
5.1.2 Generation of Exception Vector Offsets for Vectored Interrupts	
5.2 Exceptions	
5.2.1 Exception Vector Locations	
5.2.2 General Exception Processing	
MIDS22TM Architecture For Drogrammere Volume III. Devision 2.00	
ivir 352 ···· Architecture For Programmers volume III, Kevision 2.00	I

	5.2.3 EJTAG Debug Exception	36
	5.2.4 Reset Exception	36
	5.2.5 Soft Reset Exception	37
	5.2.6 Non Maskable Interrupt (NMI) Exception	38
	5.2.7 Machine Check Exception	39
	5.2.8 Address Error Exception	39
	5.2.9 TLB Refill Exception	40
	5.2.10 TLB Invalid Exception	40
	5.2.11 TLB Modified Exception	41
	5.2.12 Cache Error Exception	42
	5.2.13 Bus Error Exception	42
	5.2.14 Integer Overflow Exception	43
	5.2.15 Trap Exception	43
	5.2.16 System Call Exception	43
	5.2.17 Breakpoint Exception	43
	5.2.18 Reserved Instruction Exception	44
	5.2.19 Coprocessor Unusable Exception	44
	5.2.13 Coprocessor Character Encoption	45
	5.2.25 Flouring Four Exception	45
	5.2.21 Coprocessor 2 Exception	45
	5.2.22 Watch Exception	
	5.2.25 Interrupt Exception	-0
Chapter	6 GPR Shadow Registers	47
6.1	Introduction to Shadow Sets	47
6.2	Support Instructions	48
Chapter	7 CP0 Hazards	49
7.1	Introduction	49
7.2	Types of Hazards	49
	7.2.1 Execution Hazards	49
	7.2.2 Instruction Hazards	50
7.3	Hazard Clearing Instructions	51
	7.3.1 Instruction Encoding	51
Chantar	9 Conresson () Pagisters	52
	Coprocessor O Registers	52
0.1	Notation	55
0.2 0.2	Notation.	57
8.3	Den dem Denister (CPO Register 1, Select 0).	51
8.4	Random Register (CPO Register 1, Select 0)	50
8.5	EntryLo0, EntryLo1 (CP0 Registers 2 and 3, Select 0)	39
8.6	Context Register (CP0 Register 4, Select 0)	63
8.7	PageMask Register (CP0 Register 5, Select 0)	64
8.8	PageGrain Register (CP0 Register 5, Select 1)	66
8.9	Wired Register (CP0 Register 6, Select 0)	68
8.10	) HWREna Register (CP0 Register 7, Select 0)	69
8.1	BadVAddr Register (CP0 Register 8, Select 0)	70
8.12	2 Count Register (CP0 Register 9, Select 0)	71
8.13	3 Reserved for Implementations (CP0 Register 9, Selects 6 and 7)	71
8.14	4 EntryHi Register (CP0 Register 10, Select 0)	72
8.15	5 Compare Register (CP0 Register 11, Select 0)	74
8.16	6 Reserved for Implementations (CP0 Register 11, Selects 6 and 7)	74
8.17	7 Status Register (CP Register 12, Select 0)	75
8.18	8 IntCtl Register (CP0 Register 12, Select 1)	82
8.19	9 SRSCtl Register (CP0 Register 12, Select 2)	84
8.20	) SRSMap Register (CP0 Register 12, Select 3)	86
8.21	1 Cause Register (CP0 Register 13, Select 0)	87

8.22 Exception Program Counter (CP0 Register 14, Select 0)	
8.22.1 Special Handling of the EPC Register in Processors That Implement the MIPS16e ASE	
8.23 Processor Identification (CP0 Register 15, Select 0)	
8.24 EBase Register (CP0 Register 15, Select 1)	
8.25 Configuration Register (CP0 Register 16, Select 0)	
8.26 Configuration Register 1 (CP0 Register 16, Select 1)	
8.27 Configuration Register 2 (CP0 Register 16, Select 2)	101
8.28 Configuration Register 3 (CP0 Register 16, Select 3)	104
8.29 Reserved for Implementations (CP0 Register 16, Selects 6 and 7)	106
8.30 Load Linked Address (CP0 Register 17, Select 0)	107
8.31 WatchLo Register (CP0 Register 18)	108
8.32 WatchHi Register (CP0 Register 19)	110
8.33 Reserved for Implementations (CP0 Register 22, all Select values)	112
8.34 Debug Register (CP0 Register 23)	113
8.35 DEPC Register (CP0 Register 24)	114
8.35.1 Special Handling of the DEPC Register in Processors That Implement the MIPS16e ASE	114
8.36 Performance Counter Register (CP0 Register 25)	115
8.37 ErrCtl Register (CP0 Register 26, Select 0)	118
8.38 CacheErr Register (CP0 Register 27, Select 0)	119
8.39 TagLo Register (CP0 Register 28, Select 0, 2)	120
8.40 DataLo Register (CP0 Register 28, Select 1, 3)	121
8.41 TagHi Register (CP0 Register 29, Select 0, 2)	122
8.42 DataHi Register (CP0 Register 29, Select 1, 3)	123
8.43 ErrorEPC (CP0 Register 30, Select 0)	124
8.43.1 Special Handling of the ErrorEPC Register in Processors That Implement the MIPS16e ASE	124
8.44 DESAVE Register (CP0 Register 31)	125
Appendix A Alternative MMU Organizations	127
A.1 Fixed Mapping MMU	127
A.1.1 Fixed Address Translation	127
A.1.2 Cacheability Attributes	130
A.1.3 Changes to the CP0 Register Interface	131
A.2 Block Address Translation	131
A.2.1 BAT Organization	131
A.2.2 Address Translation	132
A.2.3 Changes to the CP0 Register Interface	133
Appendix B Revision History	135

# List of Figures

Figure 4-1: Virtual Address Space	12
Figure 4-2: References as a Function of Operating Mode	14
Figure 4-3: Contents of a TLB Entry	17
Figure 5-1: Interrupt Generation for Vectored Interrupt Mode	28
Figure 5-2: Interrupt Generation for External Interrupt Controller Interrupt Mode	30
Figure 8-1: Index Register Format	57
Figure 8-2: Random Register Format	58
Figure 8-3: EntryLo0, EntryLo1 Register Format in Release 1 of the Architecture	59
Figure 8-4: EntryLo0, EntryLo1 Register Format in Release 2 of the Architecture	60
Figure 8-5: Context Register Format	63
Figure 8-6: PageMask Register Format	64
Figure 8-7: PageGrain Register Format	66
Figure 8-8: Wired And Random Entries In The TLB	68
Figure 8-9: Wired Register Format	68
Figure 8-10: HWREna Register Format	69
Figure 8-11: BadVAddr Register Format	70
Figure 8-12: Count Register Format	71
Figure 8-13: EntryHi Register Format	72
Figure 8-14: Compare Register Format	74
Figure 8-15: Status Register Format	75
Figure 8-16: IntCtl Register Format	82
Figure 8-17: SRSCtl Register Format	84
Figure 8-18: SRSMap Register Format	86
Figure 8-19: Cause Register Format	87
Figure 8-20: EPC Register Format	91
Figure 8-21: PRId Register Format	92
Figure 8-22: EBase Register Format	93
Figure 8-23: Config Register Format	95
Figure 8-24: Config1 Register Format	97
Figure 8-25: Config2 Register Format	. 101
Figure 8-26: Config3 Register Format	. 104
Figure 8-27: LLAddr Register Format	. 107
Figure 8-28: WatchLo Register Format	. 108
Figure 8-29: WatchHi Register Format	. 110
Figure 8-30: Performance Counter Control Register Format	. 115
Figure 8-31: Performance Counter Counter Register Format	. 117
Figure 8-32: ErrorEPC Register Format	. 124
Figure 8-33: Memory Mapping when ERL = 0	. 129
Figure 8-34: Memory Mapping when ERL = 1	. 130
Figure 8-35: Config Register Additions	. 131
Figure 8-36: Contents of a BAT Entry	. 132

# List of Tables

Table 1-1: Symbols Used in Instruction Operation Statements	2
Table 4-1: Virtual Memory Address Spaces	13
Table 4-2: Address Space Access as a Function of Operating Mode	15
Table 4-3: Address Translation and Cache Coherency Attributes for the kseg0 and kseg1 Segments	
Table 4-4: Physical Address Generation	
Table 5-1: Interrupt Modes	
Table 5-2: Request for Interrupt Service in Interrupt Compatibility Mode	
Table 5-3: Relative Interrupt Priority for Vectored Interrupt Mode	
Table 5-4: Exception Vector Offsets for Vectored Interrupts	
Table 5-5: Exception Vector Base Addresses	
Table 5-6: Exception Vector Offsets	
Table 5-7: Exception Vectors	
Table 5-8: Value Stored in EPC, ErrorEPC, or DEPC on an Exception	
Table 6-1: Instructions Supporting Shadow Sets	
Table 7-1: Execution Hazards	49
Table 7-2: Instruction Hazards	50
Table 7-3: Hazard Clearing Instructions	
Table 8-1: Coprocessor 0 Registers in Numerical Order	53
Table 8-2: Read/Write Bit Field Notation	56
Table 8-3: Index Register Field Descriptions	57
Table 8-4: Random Register Field Descriptions	58
Table 8-5: EntryLo0, EntryLo1 Register Field Descriptions in Release 1 of the Architecture	59
Table 8-6: EntryLo0, EntryLo1 Register Field Descriptions in Release 2 of the Architecture	60
Table 8-7: EntryLo Field Widths as a Function of <i>PABITS</i>	<u>61</u>
Table 8-8: Cache Coherency Attributes	61
Table 8-9: Context Register Field Descriptions	63
Table 8-10: PageMask Register Field Descriptions	64
Table 8-11: Values for the Mask and MaskX <sup>1</sup> Fields of the PageMask Register	64
Table 8-12: PageGrain Register Field Descriptions	<u>66</u>
Table 8-13: Wired Register Field Descriptions	68
Table 8-14: HWREna Register Field Descriptions	69
Table 8-15: BadVAddr Register Field Descriptions	
Table 8-16: Count Register Field Descriptions	
Table 8-17: EntryHi Register Field Descriptions	
Table 8-18: Compare Register Field Descriptions	
Table 8-19: Status Register Field Descriptions	
Table 8-20: IntCtl Register Field Descriptions	
Table 8-21: SRSCtl Register Field Descriptions	
Table 8-22: Sources for new SRSCtl <sub>CSS</sub> on an Exception or Interrupt	
Table 8-23: SRSMap Register Field Descriptions	
Table 8-24: Cause Register Field Descriptions	
Table 8-25: Cause Register ExcCode Field	
Table 8-26: EPC Register Field Descriptions	
Table 8-27: PRId Register Field Descriptions	
Table 8-28: EBase Register Field Descriptions	
Table 8-29: Conditions Under Which EBase1512 Must Be Zero	
Table 8-30: Config Register Field Descriptions	
Table 8-31: Config1 Register Field Descriptions	
Table 8-32: Config2 Register Field Descriptions	101
Table 8-33: Config3 Register Field Descriptions	104

Table 8-34:	LLAddr Register Field Descriptions	107
Table 8-35:	WatchLo Register Field Descriptions	108
Table 8-36:	WatchHi Register Field Descriptions	110
Table 8-37:	Example Performance Counter Usage of the PerfCnt CP0 Register	115
Table 8-38:	Performance Counter Control Register Field Descriptions	116
Table 8-39:	Performance Counter Register Field Descriptions	117
Table 8-40:	ErrorEPC Register Field Descriptions	124
Table 8-41:	Physical Address Generation from Virtual Addresses	127
Table 8-42:	Config Register Field Descriptions	131
Table 8-43:	BAT Entry Assignments	132

### About This Book

The MIPS32<sup>TM</sup> Architecture For Programmers Volume III comes as a multi-volume set.

- Volume I describes conventions used throughout the document set, and provides an introduction to the MIPS32<sup>™</sup> Architecture
- Volume II provides detailed descriptions of each instruction in the MIPS32<sup>TM</sup> instruction set
- Volume III describes the MIPS32<sup>TM</sup> Privileged Resource Architecture which defines and governs the behavior of the privileged resources included in a MIPS32<sup>TM</sup> processor implementation
- Volume IV-a describes the MIPS16e<sup>TM</sup> Application-Specific Extension to the MIPS32<sup>TM</sup> Architecture
- Volume IV-b describes the MDMX<sup>TM</sup> Application-Specific Extension to the MIPS32<sup>TM</sup> Architecture and is not applicable to the MIPS32<sup>TM</sup> document set
- Volume IV-c describes the MIPS-3D<sup>TM</sup> Application-Specific Extension to the MIPS64<sup>TM</sup> Architecture and is not applicable to the MIPS32<sup>TM</sup> document set
- Volume IV-d describes the SmartMIPS<sup>TM</sup>Application-Specific Extension to the MIPS32<sup>TM</sup> Architecture

#### **1.1 Typographical Conventions**

This section describes the use of *italic*, **bold** and courier fonts in this book.

#### 1.1.1 Italic Text

- is used for *emphasis*
- is used for *bits, fields, registers*, that are important from a software perspective (for instance, address bits used by software, and programmable fields and registers), and various *floating point instruction formats*, such as *S*, *D*, and *PS*
- is used for the memory access types, such as cached and uncached

#### 1.1.2 Bold Text

- represents a term that is being defined
- is used for **bits** and **fields** that are important from a hardware perspective (for instance, **register** bits, which are not programmable but accessible only to hardware)
- is used for ranges of numbers; the range is indicated by an ellipsis. For instance, 5..1 indicates numbers 5 through 1
- is used to emphasize UNPREDICTABLE and UNDEFINED behavior, as defined below.

#### 1.1.3 Courier Text

Courier fixed-width font is used for text that is displayed on the screen, and for examples of code and instruction pseudocode.

#### **1.2 UNPREDICTABLE and UNDEFINED**

The terms **UNPREDICTABLE** and **UNDEFINED** are used throughout this book to describe the behavior of the processor in certain cases. **UNDEFINED** behavior or operations can occur only as the result of executing instructions in a privileged mode (i.e., in Kernel Mode or Debug Mode, or with the CP0 usable bit set in the Status register). Unprivileged software can never cause **UNDEFINED** behavior or operations. Conversely, both privileged and unprivileged software can cause **UNPREDICTABLE** results or operations.

#### **1.2.1 UNPREDICTABLE**

**UNPREDICTABLE** results may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. Software can never depend on results that are **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause a result to be generated or not. If a result is generated, it is **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause arbitrary exceptions.

**UNPREDICTABLE** results or operations have several implementation restrictions:

- Implementations of operations generating **UNPREDICTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode
- UNPREDICTABLE operations must not read, write, or modify the contents of memory or internal state which is inaccessible in the current processor mode. For example, UNPREDICTABLE operations executed in user mode must not access memory or internal state that is only accessible in Kernel Mode or Debug Mode or in another process
- UNPREDICTABLE operations must not halt or hang the processor

#### **1.2.2 UNDEFINED**

**UNDEFINED** operations or behavior may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. **UNDEFINED** operations or behavior may vary from nothing to creating an environment in which execution can no longer continue. **UNDEFINED** operations or behavior may cause data loss.

UNDEFINED operations or behavior has one implementation restriction:

• **UNDEFINED** operations or behavior must not cause the processor to hang (that is, enter a state from which there is no exit other than powering down the processor). The assertion of any of the reset signals must restore the processor to an operational state

#### 1.3 Special Symbols in Pseudocode Notation

In this book, algorithmic descriptions of an operation are described as pseudocode in a high-level language notation resembling Pascal. Special symbols used in the pseudocode notation are listed in Table 1-1.

Symbol	Meaning		
←	Assignment		
=, ≠	Tests for equality and inequality		
I	Bit string concatenation		
x <sup>y</sup>	A <i>y</i> -bit string formed by <i>y</i> copies of the single-bit value <i>x</i>		

#### Table 1-1 Symbols Used in Instruction Operation Statements

Symbol	Meaning		
b#n	A constant value $n$ in base $b$ . For instance 10#100 represents the decimal value 100, 2#100 represents the binary value 100 (decimal 4), and 16#100 represents the hexadecimal value 100 (decimal 256). If the "b#" prefix is omitted, the default base is 10.		
x <sub>yz</sub>	Selection of bits y through z of bit string x. Little-endian bit notation (rightmost bit is 0) is used. If y is less than z, this expression is an empty (zero length) bit string.		
+, -	2's complement or floating point arithmetic: addition, subtraction		
*,×	2's complement or floating point multiplication (both used for either)		
div	2's complement integer division		
mod	2's complement modulo		
/	Floating point division		
<	2's complement less-than comparison		
>	2's complement greater-than comparison		
≤	2's complement less-than or equal comparison		
2	2's complement greater-than or equal comparison		
nor	Bitwise logical NOR		
xor	Bitwise logical XOR		
and	Bitwise logical AND		
or	Bitwise logical OR		
GPRLEN	The length in bits (32 or 64) of the CPU general-purpose registers		
GPR[x]	CPU general-purpose register x. The content of <i>GPR[0]</i> is always zero.		
SGPR[s,x]	In Release 2 of the Architecture, multiple copies of the CPU general-purpose registers may be implemented. $SGPR[s,x]$ refers to GPR set <i>s</i> , register <i>x</i> . GPR[x] is a short-hand notation for $SGPR[SRSCtl_{CSS}, x]$ .		
FPR[x]	Floating Point operand register x		
FCC[CC]	Floating Point condition code CC. FCC[0] has the same value as COC[1].		
FPR[x]	Floating Point (Coprocessor unit 1), general register <i>x</i>		
CPR[z,x,s]	Coprocessor unit z, general register x, select s		
CP2CPR[x]	Coprocessor unit 2, general register <i>x</i>		
CCR[z,x]	Coprocessor unit <i>z</i> , control register <i>x</i>		
CP2CCR[x]	Coprocessor unit 2, control register <i>x</i>		
COC[z]	Coprocessor unit <i>z</i> condition signal		
Xlat[x]	Translation of the MIPS16e GPR number <i>x</i> into the corresponding 32-bit GPR number		
BigEndianMem	Endian mode as configured at chip reset (0 $\rightarrow$ Little-Endian, 1 $\rightarrow$ Big-Endian). Specifies the endianness of the memory interface (see LoadMemory and StoreMemory pseudocode function descriptions), and the endianness of Kernel and Supervisor mode execution.		
BigEndianCPU	The endianness for load and store instructions $(0 \rightarrow \text{Little-Endian}, 1 \rightarrow \text{Big-Endian})$ . In User mode, this endianness may be switched by setting the <i>RE</i> bit in the <i>Status</i> register. Thus, BigEndianCPU may be computed as (BigEndianMem XOR ReverseEndian).		

#### **Table 1-1 Symbols Used in Instruction Operation Statements**

Symbol	Meaning
ReverseEndian	Signal to reverse the endianness of load and store instructions. This feature is available in User mode only, and is implemented by setting the <i>RE</i> bit of the <i>Status</i> register. Thus, ReverseEndian may be computed as ( $SR_{RE}$ and User mode).
LLbit	Bit of <b>virtual</b> state used to specify operation for instructions that provide atomic read-modify-write. <i>LLbit</i> is set when a linked load occurs; it is tested and cleared by the conditional store. It is cleared, during other CPU operation, when a store to the location would no longer be atomic. In particular, it is cleared by exception return instructions.
I:, I+n:, I-n:	This occurs as a prefix to <i>Operation</i> description lines and functions as a label. It indicates the instruction time during which the pseudocode appears to "execute." Unless otherwise indicated, all effects of the current instruction appear to occur during the instruction time of the current instruction. No label is equivalent to a time label of <b>I</b> . Sometimes effects of an instruction appear to occur either earlier or later — that is, during the instruction time of another instruction. When this happens, the instruction operation is written in sections labeled with the instruction time, relative to the current instruction <b>I</b> , in which the effect of that pseudocode appears to occur. For example, an instruction may have a result that is not available until after the next instruction. Such an instruction has the portion of the instruction description that writes the result register in a section labeled <b>I+1</b> .
	The effect of pseudocode statements for the current instruction labelled <b>I+1</b> appears to occur "at the same time" as the effect of pseudocode statements labeled <b>I</b> for the following instruction. Within one pseudocode sequence, the effects of the statements take place in order. However, between sequences of statements for different instructions that occur "at the same time," there is no defined order. Programs must not depend on a particular order of evaluation between such sections.
РС	The <i>Program Counter</i> value. During the instruction time of an instruction, this is the address of the instruction word. The address of the instruction that occurs during the next instruction time is determined by assigning a value to <i>PC</i> during an instruction time. If no value is assigned to <i>PC</i> during an instruction time by any pseudocode statement, it is automatically incremented by either 2 (in the case of a 16-bit MIPS16e instruction) or 4 before the next instruction time. A taken branch assigns the target address to the <i>PC</i> during the instruction time of the instruction in the branch delay slot.
PABITS	The number of physical address bits implemented is represented by the symbol PABITS. As such, if 36 physical address bits were implemented, the size of the physical address space would be $2^{PABITS} = 2^{36}$ bytes.
Indicates whether the FPU has 32-bit or 64-bit floating point registers (FPRs). In MIPS32, the FP FPRs in which 64-bit data types are stored in even-odd pairs of FPRs. In MIPS64, the FPU has 3 in which 64-bit data types are stored in any FPR.	
FP32RegistersMode	In MIPS32 implementations, <b>FP32RegistersMode</b> is always a 0. MIPS64 implementations have a compatibility mode in which the processor references the FPRs as if it were a MIPS32 implementation. In such a case <b>FP32RegisterMode</b> is computed from the FR bit in the <i>Status</i> register. If this bit is a 0, the processor operates as if it had 32 32-bit FPRs. If this bit is a 1, the processor operates with 32 64-bit FPRs.
	The value of <b>FP32RegistersMode</b> is computed from the FR bit in the <i>Status</i> register.
InstructionInBranchD elaySlot	Indicates whether the instruction at the Program Counter address was executed in the delay slot of a branch or jump. This condition reflects the <i>dynamic</i> state of the instruction, not the <i>static</i> state. That is, the value is false if a branch or jump occurs to an instruction whose PC immediately follows a branch or jump, but which is not executed in the delay slot of a branch or jump.
SignalException(exce ption, argument)	Causes an exception to be signaled, using the exception parameter as the type of exception and the argument parameter as an exception-specific argument). Control does not return from this pseudocode function - the exception is signaled at the point of the call.

#### Table 1-1 Symbols Used in Instruction Operation Statements

#### **1.4 For More Information**

Various MIPS RISC processor manuals and additional information about MIPS products can be found at the MIPS URL:

http://www.mips.com

Comments or questions on the MIPS32<sup>TM</sup> Architecture or this document should be directed to

Director of MIPS Architecture MIPS Technologies, Inc. 1225 Charleston Road Mountain View, CA 94043

or via E-mail to architecture@mips.com.

### The MIPS32 Privileged Resource Architecture

#### 2.1 Introduction

The MIPS32 Privileged Resource Architecture (PRA) is a set of environments and capabilities on which the Instruction Set Architecture operates. The effects of some components of the PRA are user-visible, for instance, the virtual memory layout. Many other components are visible only to the operating system kernel and to systems programmers. The PRA provides the mechanisms necessary to manage the resources of the CPU: virtual memory, caches, exceptions and user contexts. This chapter describes these mechanisms.

#### 2.2 The MIPS Coprocessor Model

The MIPS ISA provides for up to 4 coprocessors. A coprocessor extends the functionality of the MIPS ISA, while sharing the instruction fetch and execution control logic of the CPU. Some coprocessors, such as the system coprocessor and the floating point unit are standard parts of the ISA, and are specified as such in the architecture documents. Coprocessors are generally optional, with one exception: CP0, the system coprocessor, is required. CP0 is the ISA interface to the Privileged Resource Architecture and provides full control of the processor state and modes.

#### 2.2.1 CP0 - The System Coprocessor

CP0 provides an abstraction of the functions necessary to support an operating system: exception handling, memory management, scheduling, and control of critical resources. The interface to CP0 is through various instructions encoded with the *COP0* opcode, including the ability to move data to and from the CP0 registers, and specific functions that modify CP0 state. The CP0 registers and the interaction with them make up much of the Privileged Resource Architecture.

#### 2.2.2 CP0 Registers

The CP0 registers provide the interface between the ISA and the PRA. The CP0 registers are described in Chapter 8.

# MIPS32 Operating Modes

The MIPS32 PRA requires two operating mode: User Mode and Kernel Mode. When operating in User Mode, the programmer has access to the CPU and FPU registers that are provided by the ISA and to a flat, uniform virtual memory address space. When operating in Kernel Mode, the system programmer has access to the full capabilities of the processor, including the ability to change virtual memory mapping, control the system environment, and context switch between processes.

In addition, the MIPS32 PRA supports the implementation of two additional modes: Supervisor Mode and EJTAG Debug Mode. Refer to the EJTAG specification for a description of Debug Mode.

In Release 2 of the Architecture, support was added for 64-bit coprocessors (and, in particular, 64-bit floating point units) with 32-bit CPUs. As such, certain floating point instructions which were previously enabled by 64-bit operations on a MIPS64 processor are now enabled by a new 64-bit floating point operations enabled.

#### 3.1 Debug Mode

For processors that implement EJTAG, the processor is operating in Debug Mode if the DM bit in the CP0 *Debug* register is a one. If the processor is running in Debug Mode, it has full access to all resources that are available to Kernel Mode operation.

#### 3.2 Kernel Mode

The processor is operating in Kernel Mode when the DM bit in the *Debug* register is a zero (if the processor implements Debug Mode), and any of the following three conditions is true:

- The KSU field in the CP0 Status register contains 2#00
- The EXL bit in the Status register is one
- The ERL bit in the Status register is one

The processor enters Kernel Mode at power-up, or as the result of an interrupt, exception, or error. The processor leaves Kernel Mode and enters User Mode or Supervisor Mode when all of the previous three conditions are false, usually as the result of an ERET instruction.

#### 3.3 Supervisor Mode

The processor is operating in Supervisor Mode (if that optional mode is implemented by the processor) when all of the following conditions are true:

- The DM bit in the *Debug* register is a zero (if the processor implements Debug Mode)
- The KSU field in the Status register contains 2#01
- The EXL and ERL bits in the Status register are both zero

#### 3.4 User Mode

The processor is operating in User Mode when all of the following conditions are true:

- The DM bit in the *Debug* register is a zero (if the processor implements Debug Mode)
- The KSU field in the Status register contains 2#10
- The EXL and ERL bits in the Status register are both zero

#### 3.5 Other Modes

#### 3.5.1 64-bit Floating Point Operations Enable

Instructions that are implemented by a 64-bit floating point unit are legal under any of the following conditions:

- In an implementation of Release 1 of the Architecture, 64-bit floating point operations are never enabled in a MIPS32 processor.
- If an implementation of Release 2 of the Architecture, 64-bit floating point operations are enabled if the F64 bit in the FIR register is a one. The processor must also implement the floating point data type.

#### 3.5.2 64-bit FPR Enable

Access to 64-bit FPRs is controlled by the FR bit in the *Status* register. If the FR bit is one, the FPRs are interpreted as 32 64-bit registers that may contain any data type. If the FR bit is zero, the FPRs are interpreted as 32 32-bit registers, any of which may contain a 32-bit data type (W, S). In this case, 64-bit data types are contained in even-odd pairs of registers.

64-bit FPRs are supported in a MIPS64 processor in Release 1 of the Architecture, or in a 64-bit floating point unit, for both MIPS32 and MIPS64 processors, in Release 2 of the Architecture.

The operation of the processor is UNPREDICTABLE under the following conditions:

- The FR bit is a zero, 64-bit operations are enabled, and a floating point instruction is executed whose datatype is L or PS.
- The FR bit is a zero and an odd register is referenced by an instruction whose datatype is 64-bits

### Virtual Memory

#### 4.1 Support in Release 1 and Release 2 of the Architecture

#### 4.1.1 Virtual Memory

In Release 1 of the Architecture, the minimum page size was 4KB, with optional support for pages as large as 256MB. In Release 2 of the Architecture, optional support for 1KB pages was added for use in specific embedded applications that require access to pages smaller than 4KB. Such usage is expected to be in conjunction with a default page size of 4KB and is not intended or suggested to replace the default 4KB page size but, rather, to augment it.

Support for 1KB pages involves the following changes:

- Addition of the *PageGrain* register. This register is also used by the SmartMIPS<sup>TM</sup> ASE specification, but bits used by Release 2 of the Architecture and the SmartMIPS ASE specification do not overlap.
- Modification of the EntryHi register to enable writes to, and use of, bits 12..11 (VPN2X).
- Modification of the *PageMask* register to enable writes to, and use of, bits 12..11 (MaskX).
- Modification of the *EntryLo0* and *EntryLo1* registers to shift the PFN field to the left by 2 bits, when 1KB page support is enabled, to create space for two lower-order physical address bits.

Support for 1KB pages is denoted by the Config3<sub>SP</sub> bit and enabled by the PageGrain<sub>ESP</sub> bit.

#### 4.2 Terminology

#### 4.2.1 Address Space

An *Address Space* is the range of all possible addresses that can be generated. There is one 32-bit Address Space in the MIPS32 Architecture.

#### 4.2.2 Segment and Segment Size

A *Segment* is a defined subset of an Address Space that has self-consistent reference and access behavior. Segments are either  $2^{29}$  or  $2^{31}$  bytes in size, depending on the specific Segment.

#### 4.2.3 Physical Address Size (PABITS)

The number of physical address bits implemented is represented by the symbol *PABITS*. As such, if 36 physical address bits were implemented, the size of the physical address space would be  $2^{PABITS} = 2^{36}$  bytes. The format of the *EntryLo0* and *EntryLo1* registers implicitly limits the physical address size to  $2^{36}$  bytes. Software may determine the value of PABITS by writing all ones to the *EntryLo0* or *EntryLo1* registers and reading the value back. Bits read as "1" from the PFN field allow software to determine the boundary between the PFN and 0 fields to calculate the value of PABITS.

#### 4.3 Virtual Address Spaces

The MIPS32 virtual address space is divided into five segments as shown in Figure 4-1.



Each Segment of an Address Space is classified as "Mapped" or "Unmapped". A "Mapped" address is one that is translated through the TLB or other address translation unit. An "Unmapped" address is one which is not translated through the TLB and which provides a window into the lowest portion of the physical address space, starting at physical address zero, and with a size corresponding to the size of the unmapped Segment.

Additionally, the kseg1 Segment is classified as "Uncached". References to this Segment bypass all levels of the cache hierarchy and allow direct access to memory without any interference from the caches.

Table 4-1 lists the same information in tabular form.

VA <sub>3129</sub>	Segment Name(s)	Address Range	Associated with Mode	Reference Legal from Mode(s)	Actual Segment Size
2#111	kseg3	16#FFFF FFFF through 16#E000 0000	Kernel	Kernel	2 <sup>29</sup> bytes
2#110	sseg ksseg	16#DFFF FFFF through 16#C000 0000	Supervisor	Supervisor Kernel	2 <sup>29</sup> bytes
2#101	kseg1	16#BFFF FFFF through 16#A000 0000	Kernel	Kernel	2 <sup>29</sup> bytes
2#100	kseg0	16#9FFF FFFF through 16#8000 0000	Kernel	Kernel	2 <sup>29</sup> bytes
2#0xx	useg suseg kuseg	16#7FFF FFFF through 16#0000 0000	User	User Supervisor Kernel	2 <sup>31</sup> bytes

Table 4-1	Virtual	Memory	Address	Spaces
				~~~~~~~~

Each Segment of an Address Space is associated with one of the three processor operating modes (User, Supervisor, or Kernel). A Segment that is associated with a particular mode is accessible if the processor is running in that or a more privileged mode. For example, a Segment associated with User Mode is accessible when the processor is running in User, Supervisor, or Kernel Modes. A Segment is not accessible if the processor is running in a less privileged mode than that associated with the Segment. For example, a Segment associated with Supervisor Mode is not accessible when the processor is running in User Mode and such a reference results in an Address Error Exception. The "Reference Legal from Mode(s)" column in Table 4-2 lists the modes from which each Segment may be legally referenced.

If a Segment has more than one name, each name denotes the mode from which the Segment is referenced. For example, the Segment name "useg" denotes a reference from user mode, while the Segment name "kuseg" denotes a reference to the same Segment from kernel mode.

Figure 4-2 shows the Address Space as seen when the processor is operating in each of the operating modes.

	Figur	e 4-2 References as	a Function of Opera	ting Mode	
User Mode	e References	Supervisor	Mode References	Kernel Mo	de References
16#FFFF FFFF		16#FFFF FFFF	Address Error	16#FFFF FFFF	Kamal Manad
			Address Error	Kseg3	Kernel Mapped
		16#E000 0000		16#E000 0000	
		TO#DEFE FFEF	Sama Manad	TO#DE.E.E. E.E.E.E.	S
		sseg	Supervisor Mapped	Ksseg	Supervisor Mapped
	Address Error	16#C000 0000		16#C000 0000	
		16#BFFF FFFF		16#BFFF FFFF	Kernel Unmapped
				kseg1	Uncached
			Address Error	16#A000 0000	
				16#9FFF FFFF	
				kseg0	Kernel Unmapped
16#8000 0000		16#8000 0000		16#8000 0000	
16#7FFF FFFF		16#7FFF FFFF		16#7FFF FFFF	
useg	User Mapped	suseg	User Mapped	kuseg	User Mapped
16#0000 0000		16#0000 0000		16#0000 0000	

#### 4.4 Compliance

A MIPS32 compliant processor must implement the following Segments:

- useg/kuseg
- kseg0
- kseg1

In addition, a MIPS32 compliant processor using the TLB-based address translation mechanism must also implement the kseg3 Segment.

#### 4.5 Access Control as a Function of Address and Operating Mode

Table 4-2 enumerates the action taken by the processor for each section of the 32-bit Address Space as a function of the operating mode of the processor. The selection of TLB Refill vector and other special-cased behavior is also listed for each reference.

		Action when Referenced from Operating Mode		
Virtual Address Range	Segment Name(s)	User Mode	Supervisor Mode	Kernel Mode
16#FFFF FFFF through 16#E000 0000	kseg3	Address Error	Address Error	Mapped See 4.8 on page 16 for special behavior when Debug <sub>DM</sub> = 1
16#DFFF FFFF through 16#C000 0000	sseg ksseg	Address Error	Mapped	Mapped
16#BFFF FFFF through 16#A000 0000	kseg1	Address Error	Address Error	Unmapped, Uncached See Section 4.6 on page 15
16#9FFF FFFF through 16#8000 0000	kseg0	Address Error	Address Error	Unmapped See Section 4.6 on page 15
16#7FFF FFFF through 16#0000 0000	useg suseg kuseg	Mapped	Mapped	Unmapped if Status <sub>ERL</sub> =1 See Section 4.7 on page 16 Mapped if Status <sub>ERL</sub> =0

 Table 4-2 Address Space Access as a Function of Operating Mode

#### 4.6 Address Translation and Cache Coherency Attributes for the kseg0 and kseg1 Segments

The kseg0 and kseg1 Unmapped Segments provide a window into the least significant 2<sup>29</sup> bytes of physical memory, and, as such, are not translated using the TLB or other address translation unit. The cache coherency attribute of the kseg0 Segment is supplied by the K0 field of the CP0 *Config* register. The cache coherency attribute for the kseg1 Segment is always Uncached. Table 4-3 describes how this transformation is done, and the source of the cache coherency attributes for each Segment.

Segment Name	Virtual Address Range	Generates Physical Address	Cache Attribute
	16#BFFF FFFF	16#1FFF FFFF	
kseg1	through	through	Uncached
	16#A000 0000	16#0000 0000	
	16#9FFF FFFF	16#1FFF FFFF	
kseg0	through	through	From K0 field of <i>Config</i> Register
	16#8000 0000	16#0000 0000	

#### Table 4-3 Address Translation and Cache Coherency Attributes for the kseg0 and kseg1 Segments

#### 4.7 Address Translation for the kuseg Segment when $Status_{ERL} = 1$

To provide support for the cache error handler, the kuseg Segment becomes an unmapped, uncached Segment, similar to the kseg1 Segment, if the ERL bit is set in the *Status* register. This allows the cache error exception code to operate uncached using GPR R0 as a base register to save other GPRs before use.

#### 4.8 Special Behavior for the kseg3 Segment when $Debug_{DM} = 1$

If EJTAG is implemented on the processor, the EJTAG block must treat the virtual address range 16#FF20 0000 through 16#FF3F FFFF, inclusive, as a special memory-mapped region in Debug Mode. A MIPS32 compliant implementation that also implements EJTAG must:

- explicitly range check the address range as given and not assume that the entire region between 16#FF20 0000 and 16#FFFF FFFF is included in the special memory-mapped region.
- not enable the special EJTAG mapping for this region in any mode other than in EJTAG Debug mode.

Even in Debug mode, normal memory rules may apply in some cases. Refer to the EJTAG specification for details on this mapping.

#### **4.9 TLB-Based Virtual Address Translation**<sup>1</sup>

This section describes the TLB-based virtual address translation mechanism. Note that sufficient TLB entries must be implemented to avoid a TLB exception loop on load and store instructions.

#### 4.9.1 Address Space Identifiers (ASID)

The TLB-based translation mechanism supports Address Space Identifiers to uniquely identify the same virtual address across different processes. The operating system assigns ASIDs to each process and the TLB keeps track of the ASID when doing address translation. In certain circumstances, the operating system may wish to associate the same virtual

<sup>&</sup>lt;sup>1</sup>Refer to Section A.1, "Fixed Mapping MMU" on page 127 and Section A.2, "Block Address Translation" on page 131 for descriptions of alternative MMU organizations

address with all processes. To address this need, the TLB includes a global (G) bit which over-rides the ASID comparison during translation.

#### 4.9.2 TLB Organization

The TLB is a fully-associative structure which is used to translate virtual addresses. Each entry contains two logical components: a comparison section and a physical translation section. The comparison section includes the virtual page number (VPN2 and, in Release 2, VPNX) (actually, the virtual page number/2 since each entry maps two physical pages) of the entry, the ASID, the G(lobal) bit and a recommended mask field which provides the ability to map different page sizes with a single entry. The physical translation section contains a pair of entries, each of which contains the physical page frame number (PFN), a valid (V) bit, a dirty (D) bit, and a cache coherency field (C), whose valid encodings are given in Table 8-8 on page 61. There are two entries in the translation section for each TLB entry because each TLB entry maps an aligned pair of virtual pages and the pair of physical translation entries corresponds to the even and odd pages of the pair.

Figure 4-3 shows the logical arrangement of a TLB entry, including the optional support added in Release 2 of the Architecture for 1KB page sizes. Light grey fields denote extensions to the right that are required to support 1KB page sizes. This extension is not present in an implementation of Release 1 of the Architecture.



#### Figure 4-3 Contents of a TLB Entry

Fields marked with this color are optional Release 2 features required to support 1KB pages

The fields of the TLB entry correspond exactly to the fields in the CP0 *PageMask*, *EntryHi*, *EntryLo0* and *EntryLo1* registers. The even page entries in the TLB (e.g., PFN0) come from *EntryLo0*. Similarly, odd page entries come from *EntryLo1*.

#### **4.9.3 TLB Initialization**

In many processor implementations, software must initialize the TLB during the power-up process. In processors that detect multiple TLB matches and signal this via a machine check assumption, software must be prepared to handle such an exception or use a TLB initialization algorithm that minimizes or eliminates the possibility of the exception.

In Release 1 of the Architecture, processor implementations could detect and report multiple TLB matches either on a TLB write (TLBWI or TLBWR instructions) or a TLB read (TLB access or TLBR or TLBP instructions). In Release 2 of the Architecture, processor implementations are limited to reporting multiple TLB matches only on TLB write, and this is also true of most implementations of Release 1 of the Architecture.

The following code example shows a TLB initialization routine which, on implementations of Release 2 of the Architecture, eliminates the possibility of reporting a machine check during TLB initialization. This example has equivalent effect on implementations of Release 1 of the Architecture which report multiple TLB exceptions only on a TLB write, and minimizes the probability of such an exception occuring on other implementations.

```
* InitTLB
 * Initialize the TLB to a power-up state, guaranteeing that all entries
 * are unique and invalid.
 * Arguments:
             = Maximum TLB index (from MMUSize field of C0_Config1)
      a0
 * Returns:
      No value
 *
  Restrictions:
      This routine must be called in unmapped space
 * Algorithm:
 *
      va = kseg0_base;
 *
      for (entry = max_TLB_index; entry >= 0, entry--) {
 *
          while (TLB_Probe_Hit(va)) {
 *
             va += Page_Size;
 *
          }
 *
          TLB_Write(entry, va, 0, 0, 0);
 *
      }
 *
 *
  Notes:
 *
         The Hazard macros used in the code below expand to the appropriate
 *
          number of SSNOPs in an implementation of Release 2 of the
 *
          Architecture, and to an ehb in an implementation of Release 2 of
 *
          the Architecture. See Chapter 7, "CPO Hazards," on page 49 for
 *
          more additional information.
 * /
InitTLB:
/*
 * Clear PageMask, EntryLo0 and EntryLo1 so that valid bits are off, PFN values
 * are zero, and the default page size is used.
 * /
                                   /* Clear out PFN and valid bits */
   mtc0
         zero, CO_EntryLoO
         zero, CO_EntryLol
   mtc0
          zero, CO_PageMask
                                   /* Clear out mask register *
   mtc0
/* Start with the base address of kseg0 for the VA part of the TLB */
                                    /* A_KOBASE == 16#8000.0000 */
   la
         t0, A_KOBASE
```

```
/*
* Write the VA candidate to EntryHi and probe the TLB to see if if is
* already there. If it is, a write to the TLB may cause a machine
 * check, so just increment the VA candidate by one page and try again.
*/
10:
                                   /* Write VA candidate */
   mtc0
        t0, C0_EntryHi
   TLBP_Write_Hazard()
                                   /* Clear EntryHi hazard (ssnop/ehb in R1/2) */
   tlbp
                                   /* Probe the TLB to check for a match */
   TLBP_Read_Hazard()
                                   /* Clear Index hazard (ssnop/ehb in R1/2) */
                           /* Read back flag to check for an entry */
/* Branch if about to duplicate an entry */
   mfc0 t1, C0_Index
   bgez t1, 10b
   addiu t0, (1<<S_EntryHiVPN2) /* Add 1 to VPN index in va */
/*
* A write of the VPN candidate will be unique, so write this entry
 * into the next index, decrement the index, and continue until the
 * index goes negative (thereby writing all TLB entries)
 */
        a0, C0_Index
                                   /* Use this as next TLB index */
   mtc0
                                   /* Clear Index hazard (ssnop/ehb in R1/2) */
   TLBW Write Hazard()
   tlbwi
                                   /* Write the TLB entry */
   bne a0, zero, 10b
                                   /* Branch if more TLB entries to do */
   addiu a0, -1
                                    /* Decrement the TLB index
 * Clear Index and EntryHi simply to leave the state constant for all
 * returns
 * /
         zero, C0_Index
   mtc0
   mtc0 zero, C0_EntryHi
                                    /* Return to caller */
   jr
         ra
   nop
```

#### 4.9.4 Address Translation

Release 2 of the Architecture introduced support for 1KB pages. For clarity in the discussion below, the following terms should be taken in the general sense to include the new Release 2 features:

Term Used Below	Release 2 Substitution	Comment
VPN2	VPN2    VPN2X	Release 2 implementations that support 1KB pages concatenate the VPN2 and VPN2X fields to form the virtual page number for a 1KB page
Mask	Mask    MaskX	Release 2 implementations that support 1KB pages concatenate the Mask and MaskX fields to form the don't care mask for 1KB pages

When an address translation is requested, the virtual page number and the current process ASID are presented to the TLB. All entries are checked simultaneously for a match, which occurs when all of the following conditions are true:

• The current process ASID (as obtained from the *EntryHi* register) matches the ASID field in the TLB entry, or the G bit is set in the TLB entry.

• The appropriate bits of the virtual page number match the corresponding bits of the VPN2 field stored within the TLB entry. The "appropriate" number of bits is determined by the Mask fields in each entry by ignoring each bit in the virtual page number and the TLB VPN2 field corresponding to those bits that are set in the Mask fields. This allows each entry of the TLB to support a different page size, as determined by the *PageMask* register at the time that the TLB entry was written. If the recommended *PageMask* register is not implemented, the TLB operation is as if the PageMask register was written with the encoding for a 4KB page.

If a TLB entry matches the address and ASID presented, the corresponding PFN, C, V, and D bits are read from the translation section of the TLB entry. Which of the two PFN entries is read is a function of the virtual address bit immediately to the right of the section masked with the Mask entry.

The valid and dirty bits determine the final success of the translation. If the valid bit is off, the entry is not valid and a TLB Invalid exception is raised. If the dirty bit is off and the reference was a store, a TLB Modified exception is raised. If there is an address match with a valid entry and no dirty exception, the PFN and the cache coherency bits are appended to the offset-within-page bits of the address to form the final physical address with attributes.

For clarity, the TLB lookup processes have been separated into two sets of pseudo code:

- One used by an implementation of Release 1 of the Architecture, or an implementation of Release 2 of the Architecture which does not include 1KB page support (as denoted by Config3<sub>SP</sub>). This instance is called the "4KB TLB Lookup".
- 2. One used by an implementation of Release 2 of the Architecture which does include 1KB page support. This instance is called the "1KB TLB Lookup".

The 4KB TLB Lookup pseudo code is as follows:

```
found \leftarrow 0
for i in 0...TLBEntries-1
    if ((TLB[i]<sub>VPN2</sub> and not (TLB[i]<sub>Mask</sub>)) = (va_{31..13} and not (TLB[i]<sub>Mask</sub>))) and
       (TLB[i]_{G} \text{ or } (TLB[i]_{ASID} = EntryHi_{ASID})) then
        # EvenOddBit selects between even and odd halves of the TLB as a function of
        # the page size in the matching TLB entry. Not all page sizes need
        \# be implemented on all processors, so the case below uses an `x' to
        # denote don't-care cases. The actual implementation would select
        # the even-odd bit in a way that is compatible with the page sizes
        # actually implemented.
        case TLB[i]<sub>Mask</sub>
           2#0000 0000 0000 0000: EvenOddBit \leftarrow 12 /* 4KB page */
           2#0000 0000 0000 0011: EvenOddBit <- 14 /* 16KB page */
           2#0000 0000 0000 11xx: EvenOddBit \leftarrow 16 /* 64KB page */
            2#0000 0000 0011 xxxx: EvenOddBit ← 18 /* 256KB page */
           2#0000 0000 11xx xxxx: EvenOddBit \leftarrow 20 /* 1MB page */
           2#0000 0011 xxxx xxxx: EvenOddBit \leftarrow 22 /* 4MB page */
           2#0000 llxx xxxx xxxx: EvenOddBit \leftarrow 24 /* 16MB page */
           2#0011 xxxx xxxx xxxx: EvenOddBit \leftarrow 26 /* 64MB page */
           2#11xx xxxx xxxx xxxx: EvenOddBit \leftarrow 28 /* 256MB page */
                           UNDEFINED
           otherwise:
        endcase
        if va_{EvenOddBit} = 0 then
           pfn \leftarrow TLB[i]_{PFN0}
           v \leftarrow TLB[i]_{v0}
           c \leftarrow TLB[i]_{C0}
           d \leftarrow TLB[i]_{D0}
        else
           pfn \leftarrow TLB[i]_{PFN1}
           v \leftarrow TLB[i]_{v1}
           c \leftarrow TLB[i]_{C1}
           d \leftarrow TLB[i]_{D1}
```

```
endif
         if v = 0 then
             SignalException(TLBInvalid, reftype)
         endif
         if (d = 0) and (reftype = store) then
             SignalException(TLBModified)
         endif
         # pfn<sub>PABITS-1-12..0</sub> corresponds to pa<sub>PABITS-1..12</sub>
        \texttt{pa} \leftarrow \texttt{pfn}_{\textit{PABITS-1-12..EvenOddBit-12}} ~||~ \texttt{va}_{\texttt{EvenOddBit-1..0}}
         found \leftarrow 1
        break
    endif
endfor
if found = 0 then
    SignalException(TLBMiss, reftype)
endif
```

The 1KB TLB Lookup pseudo code is as follows:

```
found \leftarrow 0
for i in 0...TLBEntries-1
   if ((TLB[i]_{VPN2} \text{ and not } (TLB[i]_{Mask})) = (va_{31..13} \text{ and not } (TLB[i]_{Mask}))) and
       (TLB[i]_G \text{ or } (TLB[i]_{ASID} = EntryHi_{ASID})) then
       # EvenOddBit selects between even and odd halves of the TLB as a function of
       # the page size in the matching TLB entry. Not all pages sizes need
       \# be implemented on all processors, so the case below uses an `x' to
       # denote don't-care cases. The actual implementation would select
       # the even-odd bit in a way that is compatible with the page sizes
       # actually implemented.
       case TLB[i]<sub>Mask</sub>
           2#0000 0000 0000 0000 00: EvenOddBit ← 10 /* 1KB page */
           2#0000 0000 0000 0000 11: EvenOddBit ← 12 /* 4KB page */
           2#0000 0000 0000 0011 xx: EvenOddBit \leftarrow 14 /* 16KB page */
           2#0000 0000 0000 11xx xx: EvenOddBit \leftarrow 16 /* 64KB page */
           2#0000 0000 0011 xxxx xx: EvenOddBit \leftarrow 18 /* 256KB page */
           2#0000 0000 11xx xxxx xx: EvenOddBit \leftarrow 20 /* 1MB page */
           2#0000 0011 xxxx xxx xx: EvenOddBit \leftarrow 22 /* 4MB page */
           2#0000 11xx xxxx xxx xx: EvenOddBit \leftarrow 24 /* 16MB page */
           2#0011 xxxx xxxx xxx xx: EvenOddBit \leftarrow 26 /* 64MB page */
           2#11xx xxxx xxxx xxx xx: EvenOddBit \leftarrow 28 /* 256MB page */
           otherwise: UNDEFINED
       endcase
       if va_{EvenOddBit} = 0 then
           pfn \leftarrow TLB[i]_{PFN0}
           v \leftarrow TLB[i]_{v0}
           c \leftarrow TLB[i]_{C0}
           d \leftarrow TLB[i]_{D0}
       else
           pfn \leftarrow TLB[i]_{PFN1}
           v \leftarrow TLB[i]_{v1}
           c \leftarrow TLB[i]_{C1}
           d \leftarrow TLB[i]_{D1}
       endif
       if v = 0 then
           SignalException(TLBInvalid, reftype)
       endif
       if (d = 0) and (reftype = store) then
           SignalException(TLBModified)
       endif
       \# pfn_{PABITS-1-10..0} corresponds to pa_{PABITS-1..10}
       pa \leftarrow pfn_{PABITS-1-10..EvenOddBit-10} \mid va_{EvenOddBit-1..0}
```

```
MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00
```

```
found ← 1
break
endif
endfor
if found = 0 then
SignalException(TLBMiss, reftype)
endif
```

Table 4-4 demonstrates how the physical address is generated as a function of the page size of the TLB entry that matches the virtual address. The "Even/Odd Select" column of Table 4-4 indicates which virtual address bit is used to select between the even (EntryLo0) or odd (EntryLo1) entry in the matching TLB entry. The "PA<sub>(PABITS-1)..0</sub> Generated From" columns specify how the physical address is generated from the selected PFN and the offset-in-page bits in the virtual address. In this column, PFN is the physical page number as loaded into the TLB from the *EntryLo0* or *EntryLo1* registers, and has one of two bit ranges:

PFN Range PA Range		Comment	
PFN(PABITS-1)-120	PA <sub>PABITS-112</sub>	Release 1 implementation, or Release 2 implementation without support for 1KB pages	
PFN(PABITS-1)-100	PAPABITS-110	Release 2 implementation with support for 1KB pages enabled	

		PA(PABITS-1)0 Generated From:		
Page Size	Even/Odd Select	Release 1 or Release 2 with 1KB Page Support Disabled	Release 2 with 1KB Page Support Enabled	
1K Bytes	VA <sub>10</sub>	Not Applicable	PFN <sub>(PABITS-1)-100</sub>    VA <sub>90</sub>	
4K Bytes	VA <sub>12</sub>	PFN <sub>(PABITS-1)-120</sub>    VA <sub>110</sub>	PFN <sub>(PABITS-1)-102</sub>    VA <sub>110</sub>	
16K Bytes	VA <sub>14</sub>	PFN <sub>(PABITS-1)-122</sub>    VA <sub>130</sub>	PFN <sub>(PABITS-1)-104</sub>    VA <sub>130</sub>	
64K Bytes	VA <sub>16</sub>	PFN <sub>(PABITS-1)-124</sub>    <sub>VA150</sub>	PFN <sub>(PABITS-1)-106</sub>    <sub>VA150</sub>	
256K Bytes	VA <sub>18</sub>	PFN <sub>(PABITS-1)-126</sub>    VA <sub>170</sub>	PFN <sub>(PABITS-1)-108</sub>    VA <sub>170</sub>	
1M Bytes	VA <sub>20</sub>	PFN <sub>(PABITS-1)-128</sub>    VA <sub>190</sub>	FN <sub>(PABITS-1)-1010</sub>    VA <sub>190</sub>	
4M Bytes	VA <sub>22</sub>	PFN <sub>(PABITS-1)-1210</sub>    VA <sub>210</sub>	PFN <sub>(PABITS-1)-1012</sub>    VA <sub>210</sub>	
16M Bytes	VA <sub>24</sub>	PFN <sub>(PABITS-1)-1212</sub>    VA <sub>230</sub>	PFN <sub>(PABITS-1)-1014</sub>    VA <sub>230</sub>	
64MBytes	VA <sub>26</sub>	PFN <sub>(PABITS-1)-1214</sub>    VA <sub>250</sub>	PFN <sub>(PABITS-1)-1016</sub>    VA <sub>250</sub>	
256MBytes	VA <sub>28</sub>	PFN <sub>(PABITS-1)-1216</sub>    VA <sub>270</sub>	PFN <sub>(PABITS-1)-1018</sub>    VA <sub>270</sub>	

#### **Table 4-4 Physical Address Generation**

# Interrupts and Exceptions

Release 2 of the Architecture added the following features related to the processing of Exceptions and Interrupts:

- The addition of the Coprocessor 0 *EBase* register, which allows the exception vector base address to be modified for exceptions that occur when Status<sub>BEV</sub> equals 0. The *EBase* register is required.
- The extension of the Release 1 interrupt control mechanism to include two optional interrupt modes:
  - Vectored Interrupt (VI) mode, in which the various sources of interrupts are prioritized by the processor and each interrupt is vectored directly to a dedicated handler. When combined with GPR shadow registers, introduced in the next chapter, this mode significantly reduces the number of cycles required to process an interrupt.
  - External Interrupt Controller (EIC) mode, in which the definition of the coprocessor 0 register fields associated with interrupts changes to support an external interrupt controller. This can support many more prioritized interrupts, while still providing the ability to vector an interrupt directly to a dedicated handler and take advantage of the GPR shadow registers.
- The ability to stop the *Count* register for highly power-sensitive applications in which the Count register is not used, or for reduced power mode. This change is required.
- The addition of the DI and EI instructions which provide the ability to atomically disable or enable interrupts. Both instructions are required.
- The addition of the TI and PCI bits in the *Cause* register to denote pending timer and performance counter interrupts. This change is required.

#### 5.1 Interrupts

Release 1 of the Architecture included support for two software interrupts, six hardware interrupts, and two special-purpose interrupts: timer and performance counter. The timer and performance counter interrupts were combined with hardware interrupt 5 in an implementation-dependent manner. Interrupts were handled either through the general exception vector (offset 16#180) or the special interrupt vector (16#200), based on the value of Cause<sub>IV</sub>. Software was required to prioritize interrupts as a function of the Cause<sub>IP</sub> bits in the interrupt handler prologue.

Release 2 of the Architecture adds an upward-compatible extension to the Release 1 interrupt architecture that supports vectored interrupts. In addition, Release 2 adds a new interrupt mode that supports the use of an external interrupt controller by changing the interrupt architecture.

Although a Non-Maskable Interrupt (NMI) includes "interrupt" in its name, it is more correctly described as an NMI exception because it does not affect, nor is it controlled by the processor interrupt system.

An interrupt is only taken when all of the following are true:

- A specific request for interrupt service is made, as a function of the interrupt mode, described below.
- The IE bit in the *Status* register is a one.
- The DM bit in the *Debug* register is a zero (for processors implementing EJTAG)
- The EXL and ERL bits in the Status register are both zero.

Logically, the request for interrupt service is ANDed with the IE bit of the *Status* register. The final interrupt request is then asserted only if both the EXL and ERL bits in the *Status* register are zero, and the DM bit in the *Debug* register is zero, corresponding to a non-exception, non-error, non-debug processing mode, respectively.

#### 5.1.1 Interrupt Modes

An implementation of Release 1 of the Architecture only implements interrupt compatibility mode.

An implementation of Release 2 of the Architecture may implement up to three interrupt modes:

- Interrupt compatibility mode, which acts identically to that in an implementation of Release 1 of the Architecture. This mode is required.
- Vectored Interrupt (VI) mode, which adds the ability to prioritize and vector interrupts to a handler dedicated to that interrupt, and to assign a GPR shadow set for use during interrupt processing. This mode is optional and its presence is denoted by the VInt bit in the *Config3* register.
- External Interrupt Controller (EIC) mode, which redefines the way in which interrupts are handled to provide full support for an external interrupt controller handling prioritization and vectoring of interrupts. This mode is optional and its presence is denoted by the VEIC bit in the *Config3* register.

A compatible implementation of Release 2 of the Architecture must implement interrupt compatibility mode, and may optionally implement one or both vectored interrupt modes. Inclusion of the optional modes may be done selectively in the implementation of the processor, or they may always be inculcated and be dynamically enabled based on coprocessor 0 control bits. The reset state of the processor is to interrupt compatibility mode such that an implementation of Release 2 of the Architecture is fully compatible with implementations of Release 1 of the Architecture.

Table 5-1 shows the current interrupt mode of the processor as a function of the coprocessor 0 register fields that can affect the mode.

Status <sub>BEV</sub>	Cause <sub>IV</sub>	IntCtl <sub>VS</sub>	Config3 <sub>VINT</sub>	Config3 <sub>VEIC</sub>	Interrupt Mode
1	х	х	х	x	Compatibly
х	0	х	х	х	Compatibility
x	x	=0	x	x	Compatibility
0	1	≠0	1	0	Vectored Interrupt
0	1	≠0	x	1	External Interrupt Controller
0	1	≠0	0	0	Can't happen - $IntCtl_{VS}$ can not be non-zero if neither Vectored Interrupt nor External Interrupt Controller mode is implemented.
			•		

#### Table 5-1 Interrupt Modes

"x" denotes don't care

#### 5.1.1.1 Interrupt Compatibility Mode

This is the only interrupt mode for a Release 1 processor and the default interrupt mode for a Release 2 processor. This mode is entered when a Reset exception occurs. In this mode, interrupts are non-vectored and dispatched though

exception vector offset 16#180 (if Cause<sub>IV</sub> = 0) or vector offset 16#200 (if Cause<sub>IV</sub> = 1). This mode is in effect if any of the following conditions are true:

- Cause<sub>IV</sub> = 0
- Status<sub>BEV</sub> = 1
- Int $Ctl_{VS} = 0$ , which would be the case if vectored interrupts are not implemented, or have been disabled.

The current interrupt requests are visible via the IP field in the Cause register on any read of the register (not just after an interrupt exception has occurred). Note that an interrupt request may be deasserted between the time the processor starts the interrupt exception and the time that the software interrupt handler runs. The software interrupt handler must be prepared to handle this condition by simply returning from the interrupt via ERET. A request for interrupt service is generated as shown in Table 5-2.

Interrupt Type	Interrupt Source	Interrupt Request Calculated From
Hardware Interrupt, Timer Interrupt, or Performance Counter Interrupt	HW5	$Cause_{IP7}$ and $Status_{IM7}$
	HW4	$Cause_{IP6}$ and $Status_{IM6}$
	HW3	$Cause_{IP5}$ and $Status_{IM5}$
Hardware Interrupt	HW2	$Cause_{IP4}$ and $Status_{IM4}$
	HW1	$Cause_{IP3}$ and $Status_{IM3}$
	HW0	$Cause_{IP2}$ and $Status_{IM2}$
Software Interment	SW1	$Cause_{IP1}$ and $Status_{IM1}$
Software Interrupt	SW0	$Cause_{IP0}$ and $Status_{IM0}$

 Table 5-2 Request for Interrupt Service in Interrupt Compatibility Mode

A typical software handler for interrupt compatibility mode might look as follows:

```
*
  Assumptions:
*
   - Cause<sub>TV</sub> = 1 (if it were zero, the interrupt exception would have to
*
                  be isolated from the general exception vector before getting
                  here)
   - GPRs k0 and k1 are available (no shadow register switches invoked in
                                   compatibility mode)
 *
   - The software priority is IP7..IP0 (HW5..HW0, SW1..SW0)
* Location: Offset 0x200 from exception base
*/
IVexception:
   mfc0
        k0, C0_Cause
                           /* Read Cause register for IP bits */
        k1, C0_Status
                             /* and Status register for IM bits */
   mfc0
        k0, k0, M_CauseIM /* Keep only IP bits from Cause */
   andi
                            /* and mask with IM bits */
         k0, k0, k1
   and
         k0, zero, Dismiss /* no bits set - spurious interrupt */
   beq
                           /* Find first bit set, IP7..IP0; k0 = 16..23 */
         k0, k0
   clz
   xori k0, k0, 0x17
                            /* 16..23 => 7..0 */
   sll
         k0, k0, VS
                            /* Shift to emulate software IntCtl<sub>vs</sub> */
   la
         k1, VectorBase
                            /* Get base of 8 interrupt vectors */
   addu
         k0, k0, k1
                             /* Compute target from base and offset */
   ir
         k0
                             /* Jump to specific exception routine */
```

```
nop
/*
 * Each interrupt processing routine processes a specific interrupt, analogous
 * to those reached in VI or EIC interrupt mode. Since each processing routine
 * is dedicated to a particular interrupt line, it has the context to know
 * which line was asserted. Each processing routine may need to look further
 * to determine the actual source of the interrupt if multiple interrupt requests
 * are ORed together on a single IP line. Once that task is performed, the
 * interrupt may be processed in one of two ways:
 * - Completely at interrupt level (e.g., a simply UART interrupt). The
 *
     SimpleInterrupt routine below is an example of this type.
 * - By saving sufficient state and re-enabling other interrupts. In this
 *
     case the software model determines which interrupts are disabled during
 *
     the processing of this interrupt. Typically, this is either the single
 *
    StatusIM bit that corresponds to the interrupt being processed, or some
 *
    collection of other Status_{TM} bits so that "lower" priority interrupts are
 *
     also disabled. The NestedInterrupt routine below is an example of this type.
 */
SimpleInterrupt:
/*
* Process the device interrupt here and clear the interupt request
 * at the device. In order to do this, some registers may need to be
 * saved and restored. The coprocessor 0 state is such that an ERET
 * will simply return to the interrupted code.
 */
                             /* Return to interrupted code */
   eret
NestedException:
/*
 * Nested exceptions typically require saving the EPC and Status registers,
 * any GPRs that may be modified by the nested exception routine, disabling
 * the appropriate IM bits in Status to prevent an interrupt loop, putting
 * the processor in kernel mode, and re-enabling interrupts. The sample code
 * below can not cover all nuances of this processing and is intended only
 * to demonstrate the concepts.
 */
   /* Save GPRs here, and setup software context */
   mfc0 k0, C0_EPC /* Get restart address */
                            /* Save in memory */
   SW
         k0, EPCSave
                            /* Get Status value */
   mfc0 k0, C0_Status
                            /* Save in memory */
         k0, StatusSave
   sw
         k1, ~IMbitsToClear /* Get Im bits to clear for this interrupt */
   li
                              /* this must include at least the IM bit */
                              /* % (1,1) = 1 for the current interrupt, and may include */
                              /* others */
   and
          k0, k0, k1
                                 /* Clear bits in copy of Status */
   ins
         k0, zero, S_StatusEXL, (W_StatusKSU+W_StatusERL+W_StatusEXL)
                                 /* Clear KSU, ERL, EXL bits in k0 */
   mtc0 k0, C0_Status
                                 /* Modify mask, switch to kernel mode, */
                                 /* re-enable interrupts */
   /*
    * Process interrupt here, including clearing device interrupt.
    * In some environments this may be done with a thread running in
    * kernel or user mode. Such an environment is well beyond the scope of
    * this example.
```

```
* /
To complete interrupt processing, the saved values must be restored
and the original interrupted code restarted.
di
                          /* Disable interrupts - may not be required */
lw
      k0, StatusSave
                          /* Get saved Status (including EXL set) */
                          /*
lw
      k1, EPCSave
                               and EPC */
mtc0 k0, C0_Status
                          /* Restore the original value */
mtc0 k1, CO_EPC
                          /* and EPC */
/* Restore GPRs and software state */
                          /* Dismiss the interrupt */
eret
```

#### 5.1.1.2 Vectored Interrupt Mode

Vectored Interrupt mode builds on the interrupt compatibility mode by adding a priority encoder to prioritize pending interrupts and to generate a vector with which each interrupt can be directed to a dedicated handler routine. This mode also allows each interrupt to be mapped to a GPR shadow set for use by the interrupt handler. Vectored Interrupt mode is in effect if all of the following conditions are true:

- Config $3_{VInt} = 1$
- Config $3_{VEIC} = 0$
- IntCtl<sub>VS</sub>  $\neq 0$
- Cause<sub>IV</sub> = 1
- Status<sub>BEV</sub> = 0

In VI interrupt mode, the six hardware interrupts are interpreted as individual hardware interrupt requests. The timer and performance counter interrupts are combined in an implementation-dependent way with the hardware interrupts (with the interrupt with which they are combined indicated by  $IntCtl_{IPTI}$  and  $IntCtl_{IPPCI}$ , respectively) to provide the appropriate relative priority of these interrupts with that of the hardware interrupts. The processor interrupt logic ANDs each of the Cause<sub>IP</sub> bits with the corresponding Status<sub>IM</sub> bits. If any of these values is 1, and if interrupts are enabled (Status<sub>IE</sub> = 1, Status<sub>EXL</sub> = 0, and Status<sub>ERL</sub> = 0), an interrupt is signaled and a priority encoder scans the values in the order shown in Table 5-3.

Relative Priority	Interrupt Type	Interrupt Source	Interrupt Request Calculated From	Vector Number Generated by Priority Encoder
Highest Priority		HW5	$Cause_{IP7}$ and $Status_{IM7}$	7
		HW4	$Cause_{IP6}$ and $Status_{IM6}$	6
	Handwana	HW3	$Cause_{IP5}$ and $Status_{IM5}$	5
HW2 Cause <sub>IP4</sub>		Cause $_{IP4}$ and $Status_{IM4}$	4	
	HW1 Cause <sub>IP3</sub> and Status <sub>IM3</sub>		3	
		HW0	$Cause_{IP2}$ and $Status_{IM2}$	2
	S - 6	SW1	$Cause_{IP1}$ and $Status_{IM1}$	1
Lowest Priority SW0 Cause <sub>IP</sub>		$Cause_{IP0}$ and $Status_{IM0}$	0	

Table 5-3 Relative Interrupt Priority for Vectored Interrupt Mode

The priority order places a relative priority on each hardware interrupt and places the software interrupts at a priority lower than all hardware interrupts. When the priority encoder finds the highest priority pending interrupt, it outputs an encoded vector number that is used in the calculation of the handler for that interrupt, as described below. This is shown pictorially in Figure 5-1.



Note that an interrupt request may be deasserted between the time the processor detects the interrupt request and the time that the software interrupt handler runs. The software interrupt handler must be prepared to handle this condition by simply returning from the interrupt via ERET.

A typical software handler for vectored interrupt mode bypasses the entire sequence of code following the IVexception label shown for the compatibility mode handler above. Instead, the hardware performs the prioritization, dispatching directly to the interrupt processing routine. Unlike the compatibility mode examples, a vectored interrupt handler may take advantage of a dedicated GPR shadow set to avoid saving any registers. As such, the SimpleInterrupt code shown above need not save the GPRs.

A nested interrupt is similar to that shown for compatibility mode, but may also take advantage of running the nested exception routine in the GPR shadow set dedicated to the interrupt or in another shadow set. Such a routine might look as follows:

```
NestedException:
/*
 * Nested exceptions typically require saving the EPC, Status and SRSCtl registers,
 * setting up the appropriate GPR shadow set for the routine, disabling
 * the appropriate IM bits in Status to prevent an interrupt loop, putting
  the processor in kernel mode, and re-enabling interrupts. The sample code
  below can not cover all nuances of this processing and is intended only
 *
  to demonstrate the concepts.
 * /
   /* Use the current GPR shadow set, and setup software context */
   mfc0 k0, C0_EPC
                            /* Get restart address */
         k0, EPCSave
                              /* Save in memory */
   sw
```

Copyright © 2001-2003 MIPS Technologies Inc. All rights reserved.
```
mfc0
         k0, C0_Status
                              /* Get Status value */
         k0, StatusSave
                              /* Save in memory */
   sw
  mfc0
         k0, C0 SRSCtl
                              /* Save SRSCtl if changing shadow sets */
         k0, SRSCtlSave
   sw
   li
         k1, ~IMbitsToClear /* Get Im bits to clear for this interrupt */
                              /*
                                 this must include at least the IM bit */
                              /*
                                   for the current interrupt, and may include */
                              /*
                                   others */
  and
         k0, k0, k1
                                 /* Clear bits in copy of Status */
   /* If switching shadow sets, write new value to {\tt SRSCtl}_{\tt PSS} here */
         k0, zero, S_StatusEXL, (W_StatusKSU+W_StatusERL+W_StatusEXL)
  ins
                                 /* Clear KSU, ERL, EXL bits in k0 */
                                 /* Modify mask, switch to kernel mode, */
  mtc0
         k0, C0_Status
                                 /*
                                      re-enable interrupts */
   /*
    * If switching shadow sets, clear only KSU above, write target
    * address to EPC, and do execute an eret to clear EXL, switch
    * shadow sets, and jump to routine
    */
   /* Process interrupt here, including clearing device interrupt */
/*
* To complete interrupt processing, the saved values must be restored
* and the original interrupted code restarted.
*/
   di
                              /* Disable interrupts - may not be required */
   lw
                            /* Get saved Status (including EXL set) */
         k0, StatusSave
                            /*
  lw
         k1, EPCSave
                                   and EPC */
                            /* Restore the original value */
  mtc0
         k0, C0_Status
                             /* Get saved SRSCtl */
         k0, SRSCtlSave
   lw
                             /* and EPC */
         k1, C0 EPC
  mtc0
                             /* Restore shadow sets */
         k0, C0_SRSCtl
  mtc0
                              /* Clear hazard */
   ehb
                              /* Dismiss the interrupt */
   eret
```

#### 5.1.1.3 External Interrupt Controller Mode

External Interrupt Controller Mode redefines the way that the processor interrupt logic is configured to provide support for an external interrupt controller. The interrupt controller is responsible for prioritizing all interrupts, including hardware, software, timer, and performance counter interrupts, and directly supplying to the processor the vector number of the highest priority interrupt. EIC interrupt mode is in effect if all of the following conditions are true:

- Config3<sub>VEIC</sub> = 1
- IntCtl<sub>VS</sub>  $\neq 0$
- Cause<sub>IV</sub> = 1
- Status<sub>BEV</sub> = 0

In EIC interrupt mode, the processor sends the state of the software interrupt requests ( $Cause_{IP1..IP0}$ ), the timer interrupt request ( $Cause_{TI}$ ), and the performance counter interrupt request ( $Cause_{PCI}$ ) to the external interrupt controller, where it prioritizes these interrupts in a system-dependent way with other hardware interrupts. The interrupt controller can be a hard-wired logic block, or it can be configurable based on control and status registers. This allows the interrupt controller to be more specific or more general as a function of the system environment and needs.

The external interrupt controller prioritizes its interrupt requests and produces the vector number of the highest priority interrupt to be serviced. The vector number, called the Requested Interrupt Priority Level (RIPL), is a 6-bit encoded

value in the range 0..63, inclusive. A value of 0 indicates that no interrupt requests are pending. The values 1..63 represent the lowest (1) to highest (63) RIPL for the interrupt to be serviced. The interrupt controller passes this value on the 6 hardware interrupt line, which are treated as an encoded value in EIC interrupt mode.

Status<sub>IPL</sub> (which overlays  $Status_{IM7..IM2}$ ) is interpreted as the Interrupt Priority Level (IPL) at which the processor is currently operating (with a value of zero indicating that no interrupt is currently being serviced). When the interrupt controller requests service for an interrupt, the processor compares RIPL with  $Status_{IPL}$  to determine if the requested interrupt has higher priority than the current IPL. If RIPL is strictly greater than  $Status_{IPL}$ , and interrupts are enabled ( $Status_{IE} = 1$ ,  $Status_{EXL} = 0$ , and  $Status_{ERL} = 0$ ) an interrupt request is signaled to the pipeline. When the processor starts the interrupt exception, it loads RIPL into  $Cause_{RIPL}$  (which overlays  $Cause_{IP7..IP2}$ ) and signals the external interrupt controller to notify it that the request is being serviced. The interrupt exception uses the value of  $Cause_{RIPL}$  as the vector number. Because  $Cause_{RIPL}$  is only loaded by the processor when an interrupt exception is signaled, it is available to software during interrupt processing.

In EIC interrupt mode, the external interrupt controller is also responsible for supplying the GPR shadow set number to use when servicing the interrupt. As such, the *SRSMap* register is not used in this mode, and the mapping of the vectored interrupt to a GPR shadow set is done by programming (or designing) the interrupt controller to provide the correct GPR shadow set number when an interrupt is requested. When the processor loads an interrupt request into Cause<sub>RIPL</sub>, it also loads the GPR shadow set number into SRSCtl<sub>EICSS</sub>, which is copied to SRSCtl<sub>CSS</sub> when the interrupt is serviced.

The operation of EIC interrupt mode is shown pictorially in Figure 5-2.

A typical software handler for EIC interrupt mode bypasses the entire sequence of code following the IV exception



Figure 5-2 Interrupt Generation for External Interrupt Controller Interrupt ModeEncodeLatchCompareGenerate

label shown for the compatibility mode handler above. Instead, the hardware performs the prioritization, dispatching directly to the interrupt processing routine. Unlike the compatibility mode examples, an EIC interrupt handler may take advantage of a dedicated GPR shadow set to avoid saving any registers. As such, the SimpleInterrupt code shown above need not save the GPRs.

A nested interrupt is similar to that shown for compatibility mode, but may also take advantage of running the nested exception routine in the GPR shadow set dedicated to the interrupt or in another shadow set. It also need only copy Cause<sub>RIPL</sub> to Status<sub>IPL</sub> to prevent lower priority interrupts from interrupting the handler. Such a routine might look as follows:

```
NestedException:
```

```
/*
* Nested exceptions typically require saving the EPC, Status, and SRSCtl registers,
* setting up the appropriate GPR shadow set for the routine, disabling
 * the appropriate IM bits in Status to prevent an interrupt loop, putting
 * the processor in kernel mode, and re-enabling interrupts. The sample code
 * below can not cover all nuances of this processing and is intended only
 * to demonstrate the concepts.
 */
   /* Use the current GPR shadow set, and setup software context */
   mfc0 k1, C0_Cause /* Read Cause to get RIPL value */
mfc0 k0, C0_EPC /* Get restart address */
   srl k1, k1, S_CauseRIPL /* Right justify RIPL field */
         k0, EPCSave /* Save in memory */
   sw
         k0, C0_Status /* Get Status value */
k0, StatusSave /* Save in memory */
   mfc0 k0, C0_Status
   SW
   ins k0, k1, S_StatusIPL, 6 /* Set IPL to RIPL in copy of Status */
   mfc0 k1, C0_SRSCt1 /* Save SRSCt1 if changing shadow sets */
         k1, SRSCtlSave
   sw
   /* If switching shadow sets, write new value to {\rm SRSCtl}_{\rm PSS} here */
   ins k0, zero, S_StatusEXL, (W_StatusKSU+W_StatusERL+W_StatusEXL)
                                  /* Clear KSU, ERL, EXL bits in k0 */
   mtc0 k0, C0_Status
                                  /* Modify IPL, switch to kernel mode, */
                                  /* re-enable interrupts */
   /*
    * If switching shadow sets, clear only KSU above, write target
    * address to EPC, and do execute an eret to clear EXL, switch
    * shadow sets, and jump to routine
    */
   /* Process interrupt here, including clearing device interrupt */
/*
 * The interrupt completion code is identical to that shown for VI mode above.
 */
```

#### 5.1.2 Generation of Exception Vector Offsets for Vectored Interrupts

For vectored interrupts (in either VI or EIC interrupt mode), a vector number is produced by the interrupt control logic. This number is combined with  $IntCtl_{VS}$  to create the interrupt offset, which is added to 16#200 to create the exception vector offset. For VI interrupt mode, the vector number is in the range 0..7, inclusive. For EIC interrupt mode, the vector number is in the range 1..63, inclusive (0 being the encoding for "no interrupt"). The  $IntCtl_{VS}$  field specifies the spacing between vector locations. If this value is zero (the default reset state), the vector spacing is zero and the processor reverts to Interrupt Compatibility Mode. A non-zero value enables vectored interrupts, and Table 5-4 shows the exception vector offset for a representative subset of the vector numbers and values of the  $IntCtl_{VS}$  field.

	Value of IntCtl <sub>VS</sub> Field				
Vector Number	2#00001	2#00010	2#00100	2#01000	2#10000
0	16#0200	16#0200	16#0200	16#0200	16#0200
1	16#0220	16#0240	16#0280	16#0300	16#0400
2	16#0240	16#0280	16#0300	16#0400	16#0600
3	16#0260	16#02C0	16#0380	16#0500	16#0800
4	16#0280	16#0300	16#0400	16#0600	16#0A00
5	16#02A0	16#0340	16#0480	16#0700	16#0C00
6	16#02C0	16#0380	16#0500	16#0800	16#0E00
7	16#02E0	16#03C0	16#0580	16#0900	16#1000
		• •			
61	16#09A0	16#1140	16#2080	16#3F00	16#7C00
62	16#09C0	16#1180	16#2100	16#4000	16#7E00
63	16#09E0	16#11C0	16#2180	16#4100	16#8000

Table 5-4 Exceptio	n Vector Offsets fo	r Vectored Interrupts
--------------------	---------------------	-----------------------

The general equation for the exception vector offset for a vectored interrupt is:

vectorOffset  $\leftarrow$  16#200 + (vectorNumber × (IntCtl<sub>VS</sub> || 2#00000))

## 5.2 Exceptions

Normal execution of instructions may be interrupted when an exception occurs. Such events can be generated as a by-product of instruction execution (e.g., an integer overflow caused by an add instruction or a TLB miss caused by a load instruction), or by an event not directly related to instruction execution (e.g., an external interrupt). When an exception occurs, the processor stops processing instructions, saves sufficient state to resume the interrupted instruction stream, enters Kernel Mode, and starts a software exception handler. The saved state and the address of the software exception handler are a function of both the type of exception, and the current state of the processor.

## 5.2.1 Exception Vector Locations

The Reset, Soft Reset, and NMI exceptions are always vectored to location 16#BFC0.0000. EJTAG Debug exceptions are vectored to location 16#BFC0.0480, or to location 16#FF20.0200 if the ProbTrap bit is zero or one, respectively, in the EJTAG\_Control\_register.

Addresses for all other exceptions are a combination of a vector offset and a vector base address. In Release 1 of the architecture, the vector base address was fixed. In Release 2 of the architecture, software is allowed to specify the vector base address via the *EBase* register for exceptions that occur when Status<sub>BEV</sub> equals 0. Table 5-5 gives the vector base address as a function of the exception and whether the BEV bit is set in the *Status* register. Table 5-6 gives the offsets from the vector base address as a function of the exception. Note that the IV bit in the *Cause* register causes Interrupts

to use a dedicated exception vector offset, rather than the general exception vector. For implementations of Release 2 of the Architecture, Table 5-4 gives the offset from the base address in the case where  $Status_{BEV} = 0$  and  $Cause_{IV} = 1$ . For implementations of Release 1 of the architecture in which  $Cause_{IV} = 1$ , the vector offset is as if  $IntCtl_{VS}$  were 0.

Table 5-7 combines these two tables into one that contains all possible vector addresses as a function of the state that can affect the vector selection. To avoid complexity in the table, the vector address value assumes that the *EBase* register, as implemented in Release 2 devices, is not changed from its reset state and that  $IntCtl_{VS}$  is 0.

In Release 2 of the Architecture, software must guarantee that EBase<sub>15..12</sub> contains zeros in all bit positions less than or equal to the most significant bit in the vector offset. This situation can only occur when a vector offset greater than 16#FFF is generated when an interrupt occurs with VI or EIC interrupt mode enabled. The operation of the processor is **UNDEFINED** if this condition is not met.

	Statı	s <sub>BEV</sub>	
Exception	0	1	
Reset, Soft Reset, NMI	16#BFC0.0000		
EJTAG Debug (with ProbEn = 0 in the EJTAG_Control_register)	16#BFC0.0480		
EJTAG Debug (with ProbEn = 1 in the EJTAG_Control_register)	16#FF2	0.0200	
	For Release 1 of the architecture:		
	16#A000.0000	16#BFC0.0300	
Casha Freeze	For Release 2 of the architecture:		
Cache Error	EBase <sub>3130</sub>    1    EBase <sub>2812</sub>    16#000		
	Note that EBase <sub>3130</sub> have the fixed value 2#10		
	For Release 1 of the architecture:		
	16#8000.0000		
Other	For Release 2 of the architecture: 16#BFC0.0200		
	EBase <sub>3112</sub>    16#000		
	Note that $EBase_{3130}$ have the fixed value $2#10$		

#### **Table 5-6 Exception Vector Offsets**

Exception	Vector Offset
TLB Refill, EXL = 0	16#000
Cache error	16#100
General Exception	16#180
Interrupt, Cause <sub>IV</sub> = 1	16#200 (In Release 2 implementations, this is the base of the vectored interrupt table when Status <sub>BEV</sub> = 0)
Reset, Soft Reset, NMI	None (Uses Reset Base Address)

					Vector
Exception	Status <sub>BEV</sub>	Status <sub>EXL</sub>	Cause <sub>IV</sub>	EJTAG ProbEn	For Release 2 Implementations, assumes that EBase retains its reset state and that IntCtl <sub>VS</sub> = 0
Reset, Soft Reset, NMI	х	х	х	х	16#BFC0.0000
EJTAG Debug	х	х	х	0	16#BFC0.0480
EJTAG Debug	х	х	х	1	16#FF20.0200
TLB Refill	0	0	х	х	16#8000.0000
TLB Refill	0	1	х	х	16#8000.0180
TLB Refill	1	0	х	х	16#BFC0.0200
TLB Refill	1	1	х	х	16#BFC0.0380
Cache Error	0	х	х	х	16#A000.0100
Cache Error	1	х	х	х	16#BFC0.0300
Interrupt	0	0	0	х	16#8000.0180
Interrupt	0	0	1	х	16#8000.0200
Interrupt	1	0	0	х	16#BFC0.0380
Interrupt	1	0	1	х	16#BFC0.0400
All others	0	х	х	х	16#8000.0180
All others	1	X	х	Х	16#BFC0.0380
'x' denotes don't care					

## Table 5-7 Exception Vectors

# 5.2.2 General Exception Processing

With the exception of Reset, Soft Reset, NMI, cache error, and EJTAG Debug exceptions, which have their own special processing as described below, exceptions have the same basic processing flow:

• If the EXL bit in the *Status* register is zero, the *EPC* register is loaded with the PC at which execution will be restarted and the BD bit is set appropriately in the *Cause* register (see Table 8-24 on page 87). The value loaded into the *EPC* register is dependent on whether the processor implements the MIPS16 ASE, and whether the instruction is in the delay slot of a branch or jump which has delay slots. Table 5-8 shows the value stored in each of the CP0 PC registers, including *EPC*. For implementations of Release 2 of the Architecture if Status<sub>BEV</sub> = 0, the CSS field in the *SRSCtl* register is copied to the PSS field, and the CSS value is loaded from the appropriate source.

If the EXL bit in the *Status* register is set, the *EPC* register is not loaded and the BD bit is not changed in the *Cause* register. For implementations of Release 2 of the Architecture, the *SRSCtl* register is not changed.

MIPS16 Implemented?	In Branch/Jump Delay Slot?	Value stored in EPC/ErrorEPC/DEPC
No	No	Address of the instruction
No	Yes	Address of the branch or jump instruction (PC-4)
Yes	No	Upper 31 bits of the address of the instruction, combined with the <i>ISA Mode</i> bit
Yes	Yes	Upper 31 bits of the branch or jump instruction (PC-2 in the MIPS16 ISA Mode and PC-4 in the 32-bit ISA Mode), combined with the <i>ISA Mode</i> bit

#### Table 5-8 Value Stored in EPC, ErrorEPC, or DEPC on an Exception

- The CE, and ExcCode fields of the *Cause* registers are loaded with the values appropriate to the exception. The CE field is loaded, but not defined, for any exception type other than a coprocessor unusable exception.
- The EXL bit is set in the Status register.
- The processor is started at the exception vector.

The value loaded into EPC represents the restart address for the exception and need not be modified by exception handler software in the normal case. Software need not look at the BD bit in the Cause register unless it wishes to identify the address of the instruction that actually caused the exception.

Note that individual exception types may load additional information into other registers. This is noted in the description of each exception type below.

#### **Operation:**

```
/* If {\rm Status}_{\rm EXL} is 1, all exceptions go through the general exception vector */
/* and neither EPC nor {\tt Cause}_{\tt BD} nor <code>SRSCtl</code> are modified */
if \text{Status}_{\text{EXL}} = 1 then
    vectorOffset \leftarrow 16#180
else
    if InstructionInBranchDelaySlot then
        EPC \leftarrow restartPC/* PC of branch/jump */
        Cause_{BD} \leftarrow 1
    else
        EPC \leftarrow restartPC
                                             /* PC of instruction */
        Cause_{BD} \leftarrow 0
    endif
    /* Compute vector offsets as a function of the type of exception */
    \texttt{NewShadowSet} \leftarrow \texttt{SRSCtl}_{\texttt{ESS}}
                                    /* Assume exception, Release 2 only */
    if ExceptionType = TLBRefill then
        vectorOffset \leftarrow 16#000
    elseif (ExceptionType = Interrupt) then
        if (Cause_{TV} = 0) then
            vectorOffset \leftarrow 16#180
        else
            if (Status<sub>BEV</sub> = 1) or (IntCtl<sub>VS</sub> = 0) then
                vectorOffset \leftarrow 16#200
            else
                if Config3_{VEIC} = 1 then
                    VecNum \leftarrow Cause_{RIPL}
                    NewShadowSet \leftarrow SRSCtl<sub>EICSS</sub>
                else
```

```
\texttt{NewShadowSet} \leftarrow \texttt{SRSMap}_{\texttt{IPL}} \times_{\texttt{4+3..IPL}} \times_{\texttt{4}}
                 endif
                 vectorOffset \leftarrow 16#200 + (VecNum × (IntCtl<sub>VS</sub> || 2#00000))
            endif /* if (Status_{\rm BEV} = 1) or (IntCtl_{\rm VS} = 0) then */
        endif /* if (Cause<sub>IV</sub> = 0) then */
    endif /* elseif (ExceptionType = Interrupt) then */
    /* Update the shadow set information for an implementation of */
    /* Release 2 of the architecture */
    if (ArchitectureRevision \geq 2) and (SRSCtl<sub>HSS</sub> > 0) and (Status<sub>BEV</sub> = 0) then
        SRSCtl_{PSS} \leftarrow SRSCtl_{CSS}
        SRSCtl_{CSS} \leftarrow NewShadowSet
    endif
endif /* if Status_{EXL} = 1 then */
Cause_{CE} \leftarrow FaultingCoprocessorNumber
Cause_{ExcCode} \leftarrow ExceptionType
Status_{EXL} \leftarrow 1
/* Calculate the vector base address */
if Status_{BEV} = 1 then
    vectorBase \leftarrow 16#BFC0.0200
else
    if ArchitectureRevision \geq 2 then
        /* The fixed value of \textsc{EBase}_{31\dots30} forces the base to be in kseg0 or kseg1 */
        vectorBase \leftarrow EBase<sub>31..12</sub> || 16#000
    else
        vectorBase ← 16#8000.0000
    endif
endif
/* Exception PC is the sum of vectorBase and vectorOffset. Vector */
/* offsets > 16#FFF (vectored or EIC interrupts only), require */
/* that \mathtt{EBase}_{15\,.\,12} have zeros in each bit position less than or */
/* equal to the most significant bit position of the vector offset */
PC \leftarrow vectorBase_{31..30} \parallel (vectorBase_{29..0} + vectorOffset_{29..0})
                                  /* No carry between bits 29 and 30 */
```

## 5.2.3 EJTAG Debug Exception

An EJTAG Debug Exception occurs when one of a number of EJTAG-related conditions is met. Refer to the EJTAG Specification for details of this exception.

#### **Entry Vector Used**

16#BFC0 0480 if the ProbTrap bit is zero in the EJTAG\_Control\_register; 16#FF20 0200 if the ProbTrap bit is one.

#### 5.2.4 Reset Exception

A Reset Exception occurs when the Cold Reset signal is asserted to the processor. This exception is not maskable. When a Reset Exception occurs, the processor performs a full reset initialization, including aborting state machines, establishing critical state, and generally placing the processor in a state in which it can execute instructions from uncached, unmapped address space. On a Reset Exception, only the following registers have defined state:

• The *Random* register is initialized to the number of TLB entries - 1.

- The Wired register is initialized to zero.
- The Config, Config1, Config2, and Config3 registers are initialized with their boot state.
- The RP, BEV, TS, SR, NMI, and ERL fields of the Status register are initialized to a specified state.
- Watch register enables and Performance Counter register interrupt enables are cleared.
- The *ErrorEPC* register is loaded with the restart PC, as described in Table 5-8. Note that this value may or may not be predictable if the Reset Exception was taken as the result of power being applied to the processor because PC may not have a valid value in that case. In some implementations, the value loaded into *ErrorEPC* register may not be predictable on either a Reset or Soft Reset Exception.
- PC is loaded with 16#BFC0 0000.

#### Cause Register ExcCode Value

None

#### **Additional State Saved**

None

#### Entry Vector Used

Reset (16#BFC0 0000)

#### Operation

```
Random \leftarrow TLBEntries - 1
Wired \leftarrow 0
Config ← ConfigurationState
                                          # Suggested - see Config register description
Config_{K0} \leftarrow 2
Config1 \leftarrow ConfigurationState
Config2 \leftarrow ConfigurationState \# if implemented
Config3 ← ConfigurationState # if implemented
Status_{RP} \leftarrow 0
\text{Status}_{\text{BEV}} \leftarrow 1
\text{Status}_{\text{TS}} \leftarrow 0
\texttt{Status}_{\texttt{SR}} \ \leftarrow \ \texttt{0}
\text{Status}_{\text{NMI}} \leftarrow 0
\text{Status}_{\text{ERL}} \leftarrow 1
WatchLo[n]<sub>T</sub> \leftarrow 0
                                       # For all implemented Watch registers
WatchLo[n]_R \leftarrow 0
                                       # For all implemented Watch registers
WatchLo[n]_W \leftarrow 0
                                        # For all implemented Watch registers
PerfCnt.Control[n]_{IE} \leftarrow 0
                                        # For all implemented PerfCnt registers
if InstructionInBranchDelaySlot then
    ErrorEPC ← restartPC # PC of branch/jump
else
    ErrorEPC \leftarrow restartPC \ \# PC \ of \ instruction
endif
PC ← 16#BFC0 0000
```

#### 5.2.5 Soft Reset Exception

A Soft Reset Exception occurs when the Reset signal is asserted to the processor. This exception is not maskable. When a Soft Reset Exception occurs, the processor performs a subset of the full reset initialization. Although a Soft Reset Exception does not unnecessarily change the state of the processor, it may be forced to do so in order to place the processor in a state in which it can execute instructions from uncached, unmapped address space. Since bus, cache, or other operations may be interrupted, portions of the cache, memory, or other processor state may be inconsistent. The primary difference between the Reset and Soft Reset Exceptions is in actual use. The Reset Exception is typically used to initialize the processor on power-up, while the Soft Reset Exception is typically used to recover from a non-responsive (hung) processor. The semantic difference is provided to allow boot software to save critical coprocessor 0 or other register state to assist in debugging the potential problem. As such, the processor may reset the same state when either reset signal is asserted, but the interpretation of any state saved by software may be very different.

In addition to any hardware initialization required, the following state is established on a Soft Reset Exception:

- The RP, BEV, TS, SR, NMI, and ERL fields of the Status register are initialized to a specified state.
- Watch register enables and Performance Counter register interrupt enables are cleared.
- The *ErrorEPC* register is loaded with the restart PC, as described in Table 5-8.
- PC is loaded with 16#BFC0 0000.

#### Cause Register ExcCode Value

None

**Additional State Saved** 

None

#### **Entry Vector Used**

Reset (16#BFC0 0000)

#### Operation

```
Config_{K0} \leftarrow 2
                                          # Suggested - see Config register description
\text{Status}_{\text{RP}} \leftarrow 0
Status_{BEV} \leftarrow 1
\text{Status}_{\text{TS}} \leftarrow 0
\texttt{Status}_{\texttt{SR}} \, \leftarrow \, \texttt{1}
\text{Status}_{\text{NMI}} \leftarrow 0
\texttt{Status}_{\texttt{ERL}} \ \leftarrow \ \texttt{1}
WatchLo[n]<sub>I</sub> \leftarrow 0
                                          # For all implemented Watch registers
WatchLo[n]_R \leftarrow 0
                                          # For all implemented Watch registers
WatchLo[n]_W \leftarrow 0
                                          # For all implemented Watch registers
PerfCnt.Control[n]_{IE} \leftarrow 0
                                        # For all implemented PerfCnt registers
if InstructionInBranchDelaySlot then
    else
    ErrorEPC \leftarrow restartPC \# PC of instruction
endif
PC ← 16#BFC0 0000
```

#### 5.2.6 Non Maskable Interrupt (NMI) Exception

A non maskable interrupt exception occurs when the NMI signal is asserted to the processor.

Although described as an interrupt, it is more correctly described as an exception because it is not maskable. An NMI occurs only at instruction boundaries, so does not do any reset or other hardware initialization. The state of the cache, memory, and other processor state is consistent and all registers are preserved, with the following exceptions:

- The BEV, TS, SR, NMI, and ERL fields of the Status register are initialized to a specified state.
- The *ErrorEPC* register is loaded with restart PC, as described in Table 5-8.
- PC is loaded with 16#BFC0 0000.

#### Cause Register ExcCode Value

None

Additional State Saved

None

#### **Entry Vector Used**

Reset (16#BFC0 0000)

### Operation

## 5.2.7 Machine Check Exception

A machine check exception occurs when the processor detects an internal inconsistency.

The following conditions cause a machine check exception:

• Detection of multiple matching entries in the TLB in a TLB-based MMU.

#### Cause Register ExcCode Value

MCheck (See Table 8-25 on page 90)

#### **Additional State Saved**

Depends on the condition that caused the exception. See the descriptions above.

#### **Entry Vector Used**

General exception vector (offset 16#180)

#### 5.2.8 Address Error Exception

An address error exception occurs under the following circumstances:

- An instruction is fetched from an address that is not aligned on a word boundary.
- A load or store word instruction is executed in which the address is not aligned on a word boundary.
- A load or store halfword instruction is executed in which the address is not aligned on a halfword boundary.
- A reference is made to a kernel address space from User Mode or Supervisor Mode.
- A reference is made to a supervisor address space from User Mode.

Note that in the case of an instruction fetch that is not aligned on a word boundary, the PC is updated before the condition is detected. Therefore, both EPC and BadVAddr point at the unaligned instruction address.

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

#### Copyright © 2001-2003 MIPS Technologies Inc. All rights reserved.

#### Cause Register ExcCode Value

AdEL: Reference was a load or an instruction fetch

AdES: Reference was a store

See Table 8-25 on page 90.

## Additional State Saved

<b>Register State</b>	Value
BadVAddr	failing address
Context <sub>VPN2</sub>	UNPREDICTABLE
EntryHi <sub>VPN2</sub>	UNPREDICTABLE
EntryLo0	UNPREDICTABLE
EntryLo1	UNPREDICTABLE

#### **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.9 TLB Refill Exception

A TLB Refill exception occurs in a TLB-based MMU when no TLB entry matches a reference to a mapped address space and the EXL bit is zero in the *Status* register. Note that this is distinct from the case in which an entry matches but has the valid bit off, in which case a TLB Invalid exception occurs.

## Cause Register ExcCode Value

TLBL: Reference was a load or an instruction fetch

TLBS: Reference was a store

See Table 8-25 on page 90.

## **Additional State Saved**

<b>Register State</b>	Value
BadVAddr	failing address
Context	The BadVPN2 field contains $VA_{3113}$ of the failing address
EntryHi	The VPN2 field contains $VA_{3113}$ of the failing address; the ASID field contains the ASID of the reference that missed.
EntryLo0	UNPREDICTABLE
EntryLo1	UNPREDICTABLE

#### **Entry Vector Used**

- TLB Refill vector (offset 16#000) if Status<sub>EXL</sub> = 0 at the time of exception.
- General exception vector (offset 16#180) if  $\text{Status}_{\text{EXL}} = 1$  at the time of exception

#### 5.2.10 TLB Invalid Exception

A TLB invalid exception occurs when a TLB entry matches a reference to a mapped address space, but the matched entry has the valid bit off.

Note that the condition in which no TLB entry matches a reference to a mapped address space and the EXL bit is one in the *Status* register is indistinguishable from a TLB Invalid Exception in the sense that both use the general exception vector and supply an ExcCode value of TLBL or TLBS. The only way to distinguish these two cases is by probing the TLB for a matching entry (using TLBP).

### *Cause* Register ExcCode Value

TLBL: Reference was a load or an instruction fetch

TLBS: Reference was a store

See Table 8-24 on page 87.

## **Additional State Saved**

<b>Register State</b>	Value
BadVAddr	failing address
Context	The BadVPN2 field contains $VA_{3113}$ of the failing address
EntryHi	The VPN2 field contains $VA_{3113}$ of the failing address; the ASID field contains the ASID of the reference that missed.
EntryLo0	UNPREDICTABLE
EntryLo1	UNPREDICTABLE

#### **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.11 TLB Modified Exception

A TLB modified exception occurs on a *store* reference to a mapped address when the matching TLB entry is valid, but the entry's D bit is zero, indicating that the page is not writable.

## Cause Register ExcCode Value

Mod (See Table 8-24 on page 87)

#### **Additional State Saved**

<b>Register State</b>	Value
BadVAddr	failing address
Context	The BadVPN2 field contains $VA_{3113}$ of the failing address
EntryHi	The VPN2 field contains $VA_{3113}$ of the failing address; the ASID field contains the ASID of the reference that missed.
EntryLo0	UNPREDICTABLE
EntryLo1	UNPREDICTABLE

#### **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.12 Cache Error Exception

A cache error exception occurs when an instruction or data reference detects a cache tag or data error, or a parity or ECC error is detected on the system bus when a cache miss occurs. This exception is not maskable. Because the error was in a cache, the exception vector is to an unmapped, uncached address.

#### Cause Register ExcCode Value

N/A

**Additional State Saved** 

Register State	Value
CacheErr	Error state
ErrorEPC	Restart PC

## **Entry Vector Used**

Cache error vector (offset 16#100)

## Operation

```
CacheErr \leftarrow ErrorState

Status<sub>ERL</sub> \leftarrow 1

if InstructionInBranchDelaySlot then

ErrorEPC \leftarrow restartPC # PC of branch/jump

else

ErrorEPC \leftarrow restartPC # PC of instruction

endif

if Status<sub>BEV</sub> = 1 then

PC \leftarrow 16#BFC0 0200 + 16#100

else

PC \leftarrow 16#A000 0000 + 16#100

endif
```

## 5.2.13 Bus Error Exception

A bus error occurs when an instruction, data, or prefetch access makes a bus request (due to a cache miss or an uncacheable reference) and that request is terminated in an error. Note that parity errors detected during bus transactions are reported as cache error exceptions, not bus error exceptions.

#### Cause Register ExcCode Value

IBE: Error on an instruction reference

DBE: Error on a data reference

See Table 8-25 on page 90.

#### **Additional State Saved**

None

#### **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.14 Integer Overflow Exception

An integer overflow exception occurs when selected integer instructions result in a 2's complement overflow.

## Cause Register ExcCode Value

Ov (See Table 8-25 on page 90)

## **Additional State Saved**

None

#### **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.15 Trap Exception

A trap exception occurs when a trap instruction results in a TRUE value.

Cause Register ExcCode Value

Tr (See Table 8-25 on page 90)

**Additional State Saved** 

None

# Entry Vector Used General exception vector (offset 16#180)

## 5.2.16 System Call Exception

A system call exception occurs when a SYSCALL instruction is executed.

## Cause Register ExcCode Value

Sys (See Table 8-24 on page 87)

## Additional State Saved

None

## **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.17 Breakpoint Exception

A breakpoint exception occurs when a BREAK instruction is executed.

## Cause Register ExcCode Value

Bp (See Table 8-25 on page 90)

## **Additional State Saved**

None

## **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.18 Reserved Instruction Exception

A Reserved Instruction Exception occurs if any of the following conditions is true:

- An instruction was executed that specifies an encoding of the opcode field that is flagged with "\*" (reserved), "β" (higher-order ISA), or an unimplemented "ε" (ASE).
- An instruction was executed that specifies a *SPECIAL* opcode encoding of the function field that is flagged with "\*" (reserved), or "β" (higher-order ISA).
- An instruction was executed that specifies a *REGIMM* opcode encoding of the rt field that is flagged with "\*" (reserved).
- An instruction was executed that specifies an unimplemented *SPECIAL2* opcode encoding of the function field that is flagged with an unimplemented " $\theta$ " (partner available), or an unimplemented " $\sigma$ " (EJTAG).
- An instruction was executed that specifies a *COPz* opcode encoding of the rs field that is flagged with "\*" (reserved), "β" (higher-order ISA), or an unimplemented "ε" (ASE), assuming that access to the coprocessor is allowed. If access to the coprocessor is not allowed, a Coprocessor Unusable Exception occurs instead. For the *COP1* opcode, some implementations of previous ISAs reported this case as a Floating Point Exception, setting the Unimplemented Operation bit in the Cause field of the *FCSR* register.
- An instruction was executed that specifies an unimplemented *COP0* opcode encoding of the function field when rs is *CO* that is flagged with "\*" (reserved), or an unimplemented "σ" (EJTAG), assuming that access to coprocessor 0 is allowed. If access to the coprocessor is not allowed, a Coprocessor Unusable Exception occurs instead.
- An instruction was executed that specifies a *COP1* opcode encoding of the function field that is flagged with "\*" (reserved), "β" (higher-order ISA), or an unimplemented "ε" (ASE), assuming that access to coprocessor 1 is allowed. If access to the coprocessor is not allowed, a Coprocessor Unusable Exception occurs instead. Some implementations of previous ISAs reported this case as a Floating Point Exception, setting the Unimplemented Operation bit in the Cause field of the *FCSR* register.

## Cause Register ExcCode Value

RI (See Table 8-25 on page 90)

#### **Additional State Saved**

None

## **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.19 Coprocessor Unusable Exception

A coprocessor unusable exception occurs if any of the following conditions is true:

- A COP0 or Cache instruction was executed while the processor was running in a mode other than Debug Mode or Kernel Mode, and the CU0 bit in the *Status* register was a zero
- A COP1, LWC1, SWC1, LDC1, SDC1 or MOVCI (Special opcode function field encoding) instruction was executed and the CU1 bit in the *Status* register was a zero.
- A COP2, LWC2, SWC2, LDC2, or SDC2 instruction was executed, and the CU2 bit in the Status register was a zero.
- A COP3 instruction was executed, and the CU3 bit in the *Status* register was a zero.

## Cause Register ExcCode Value

CpU (See Table 8-24 on page 87)

**Additional State Saved** 

**Register State** 

Value

Cause<sub>CE</sub> unit number of the coprocessor being referenced

## **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.20 Floating Point Exception

A floating point exception is initiated by the floating point coprocessor to signal a floating point exception.

**Register ExcCode Value** 

FPE (See Table 8-24 on page 87)

**Additional State Saved** 

<b>Register State</b>	Value
FCSR	indicates the cause of the floating point exception

## **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.21 Coprocessor 2 Exception

A coprocessor 2 exception is initiated by coprocessor 2 to signal a precise coprocessor 2 exception.

#### **Register ExcCode Value**

C2E (See Table 8-24 on page 87)

#### Additional State Saved

Defined by the coprocessor

#### **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.22 Watch Exception

The watch facility provides a software debugging vehicle by initiating a watch exception when an instruction or data reference matches the address information stored in the *WatchHi* and *WatchLo* registers. A watch exception is taken immediately if the EXL and ERL bits of the *Status* register are both zero. If either bit is a one at the time that a watch exception would normally be taken, the WP bit in the *Cause* register is set, and the exception is deferred until both the EXL and ERL bits in the Status register are zero. Software may use the WP bit in the *Cause* register to determine if the EPC register points at the instruction that caused the watch exception, or if the exception actually occurred while in kernel mode.

If the EXL or ERL bits are one in the *Status* register and a single instruction generates both a watch exception (which is deferred by the state of the EXL and ERL bits) and a lower-priority exception, the lower priority exception is taken.

Watch exceptions are never taken if the processor is executing in Debug Mode. Should a watch register match while the processor is in Debug Mode, the exception is inhibited and the WP bit is not changed.

It is implementation dependent whether a data watch exception is triggered by a prefetch or cache instruction whose address matches the Watch register address match conditions. A watch triggered by a SC instruction does so even if the store would not complete because the LLbit is zero.

## **Register ExcCode Value**

WATCH (See Table 8-24 on page 87)

## **Additional State Saved**

Register State	Value
Cause <sub>WP</sub>	indicates that the watch exception was deferred until after both Status <sub>EXL</sub> and Status <sub>ERL</sub> were zero. This bit directly causes a watch exception, so software must clear this bit as part of the exception handler to prevent a watch exception loop at the end of the current handler execution.

#### **Entry Vector Used**

General exception vector (offset 16#180)

## 5.2.23 Interrupt Exception

The interrupt exception occurs when an enabled request for interrupt service is made. See Section 5.1 on page 23 for more information.

#### **Register ExcCode Value**

Int (See Table 8-25 on page 90)

#### **Additional State Saved**

Register State	Value
Cause <sub>IP</sub>	indicates the interrupts that are pending.

#### **Entry Vector Used**

General exception vector (offset 16#180) if the IV bit in the *Cause* register is zero. Interrupt vector (offset 16#200) if the IV bit in the *Cause* register is one.

# **GPR Shadow Registers**

The capability in this chapter is targeted at removing the need to save and restore GPRs on entry to high priority interrupts or exceptions, and to provide specified processor modes with the same capability. This is done by introducing multiple copies of the GPRs, called *shadow sets*, and allowing privileged software to associate a shadow set with entry to Kernel Mode via an interrupt vector or exception. The normal GPRs are logically considered shadow set zero.

The number of GPR shadow sets is implementation dependent and may range from one (the normal GPRs) to an architectural maximum of 16. The highest number actually implemented is indicated by the  $SRSCtl_{HSS}$  field. If this field is zero, only the normal GPRs are implemented.

# 6.1 Introduction to Shadow Sets

Shadow sets are new copies of the GPRs that can be substituted for the normal GPRs on entry to Kernel Mode via an interrupt or exception. Once a shadow set is bound to a Kernel Mode entry condition, reference to GPRs work exactly as one would expect, but they are redirected to registers that are dedicated to that condition. Privileged software may need to reference all GPRs in the register file, even specific shadow registers that are not visible in the current mode. The RDPGPR and WRPGPR instructions are used for this purpose. The CSS field of the *SRSCtl* register provides the number of the current shadow register set, and the PSS field of the *SRSCtl* register provides the number of the previous shadow register set (that which was current before the last exception or interrupt occurred).

If the processor is operating in VI interrupt mode, binding of a vectored interrupt to a shadow set is done by writing to the *SRSMap* register. If the processor is operating in EIC interrupt mode, the binding of the interrupt to a specific shadow set is provided by the external interrupt controller, and is configured in an implementation-dependent way. Binding of an exception or non-vectored interrupt to a shadow set is done by writing to the ESS field of the *SRSCt1* register. When an exception or interrupt occurs, the value of SRSCt1<sub>CSS</sub> is copied to SRSCt1<sub>PSS</sub>, and SRSCt1<sub>CSS</sub> is set to the value taken from the appropriate source. On an ERET, the value of SRSCt1<sub>PSS</sub> is copied back into SRSCt1<sub>CSS</sub> to restore the shadow set of the mode to which control returns. More precisely, the rules for updating the fields in the *SRSCt1* register on an interrupt or exception are as follows:

- 1. No field in the *SRSCtl* register is updated if any of the following conditions is true. In this case, steps 2 and 3 are skipped.
  - The exception is one that sets Status<sub>ERL</sub>: NMI or cache error.
  - The exception causes entry into EJTAG Debug Mode
  - Status<sub>BEV</sub> = 1
  - Status<sub>EXL</sub> = 1
- 2. SRSCtl<sub>CSS</sub> is copied to SRSCtl<sub>PSS</sub>
- 3. SRSCtl<sub>CSS</sub> is updated from one of the following sources:
  - The appropriate field of the *SRSMap* register, based on IPL, if the exception is an interrupt, Cause<sub>IV</sub> = 1, Config $_{VEIC}$  = 0, and Config $_{VInt}$  = 1. These are the conditions for a vectored interrupt.
  - The EICSS field of the *SRSCtl* register if the exception is an interrupt,  $Cause_{IV} = 1$  and  $Config3_{VEIC} = 1$ . These are the conditions for a vectored EIC interrupt.
  - The ESS field of the *SRSCtl* register in any other case. This is the condition for a non-interrupt exception, or a non-vectored interrupt.

Similarly, the rules for updating the fields in the SRSCtl register at the end of an exception or interrupt are as follows:

- 1. No field in the SRSCtl register is updated if any of the following conditions is true. In this case, step 2 is skipped.
  - A DERET is executed
  - An ERET is executed with  $\text{Status}_{\text{ERL}} = 1$  or  $\text{Status}_{\text{BEV}} = 1$
- 2. SRSCtl<sub>PSS</sub> is copied to SRSCtl<sub>CSS</sub>

These rules have the effect of preserving the *SRSCtl* register in any case of a nested exception or one which occurs before the processor has been fully initialize (Status<sub>BEV</sub> = 1).

Privileged software may switch the current shadow set by writing a new value into SRSCtl<sub>PSS</sub>, loading EPC with a target address, and doing an ERET.

# 6.2 Support Instructions

Mnemonic Function		MIPS64 Only?
RDPGPR	Read GPR From Previous Shadow Set	No
WRPGPR	Write GPR to Shadow Set	No

#### Table 6-1 Instructions Supporting Shadow Sets

# CP0 Hazards

# 7.1 Introduction

Because resources controlled via Coprocessor 0 affect the operation of various pipeline stages of a MIPS32 processor, manipulation of these resources may produce results that are not detectable by subsequent instructions for some number of execution cycles. When no hardware interlock exists between one instruction that causes an effect that is visible to a second instruction, a *CP0 hazard* exists.

In Release 1 of the MIPS32<sup>TM</sup> Architecture, CP0 hazards were relegated to implementation-dependent cycle-based solutions, primarily based on the SSNOP instruction. Since that time, it has become clear that this is an insufficient and error-prone practice that must be addressed with a firm compact between hardware and software. As such, new instructions have been added to Release 2 of the architecture which act as explicit barriers that eliminate hazards. To the extent that it was possible to do so, the new instructions have been added in such a way that they are backward-compatible with existing MIPS processors.

# 7.2 Types of Hazards

In privileged software, there are two different types of hazards: execution hazards and instruction hazards. Both are defined below. In Table 7-1 and Table 7-2 below, the final column lists the "typical" spacing required in implementations of Release 1 of the Architecture to allow the consumer to eliminate the hazard. The "typical" value shown in these tables represent spacing that is in common use by operating systems today. An implementation of Release 1 of the Architecture which requires less spacing to clear the hazard (including one which has full hardware interlocking) should operate correctly with an operating system which uses this hazard table. An implementation of Release 1 of the Architecture which requires more spacing to clear the hazard incurs the burden of validating kernel code against the new hazard requirements.

Note that, for superscalar MIPS implementations, the number of instructions issued per cycle may be greater than one, and thus that the duration of the hazard in instructions may be greater than the duration in cycles. It is for this reason that MIPS32 Release 1 defines the SSNOP instruction to convert instruction issues to cycles in a superscalar design.

## 7.2.1 Execution Hazards

Execution hazards are those created by the execution of one instruction, and seen by the execution of another instruction. Table 7-1 lists execution hazards.

Producer	$\rightarrow$	Consumer	Hazard On	"Typical" Spacing (Cycles)
TI BWD TI BWI	, TLBWI $\rightarrow$	TLBP, TLBR	TLB entry	3
ILDWK, ILDWI		Load/store using new TLB entry	TLB entry	3
MTC0	$\rightarrow$	Load/store affected by new state	EntryHi <sub>ASID</sub> WatchHi WatchLo	3

#### **Table 7-1 Execution Hazards**

Producer	$\rightarrow$	Consumer	Hazard On	"Typical" Spacing (Cycles)
MTC0	$\rightarrow$	Coprocessor instruction execution depends on the new value of $\ensuremath{Status}_{CU}$	Status <sub>CU</sub>	4
MTC0	$\rightarrow$	ERET	Status EPC DEPC ErrorEPC	3
MTC0, EI, DI	$\rightarrow$	Interrupted Instruction	Status <sub>IE</sub>	3
MTC0	$\rightarrow$	Interrupted Instruction	Cause <sub>IP</sub>	3
MTC0	$\rightarrow$	Interrupted Instruction	Compare	3
MTC0	$\rightarrow$	CACHE	PageGrain	2
TLBR	$\rightarrow$ MFC0		EntryHi, EntryLo0, EntryLo1, PageMask	3
TLBP	$\rightarrow$	MFC0	Index	2
MTC0	$\rightarrow$	TLBR TLBWI TLBWR	EntryHi	2
MTC0	$\rightarrow$	TLBP Load or Store Instruction	EntryHi <sub>ASID</sub>	3
MTC0	$\rightarrow$	TLBWI Inc TLBWR Entr Entr		2
MTC0	$\rightarrow$	RDPGPR WRPGPR	SRSCtl <sub>PSS</sub>	2
LL	$\rightarrow$	MFC0	LLAddr	2

## Table 7-1 Execution Hazards

## 7.2.2 Instruction Hazards

Instruction hazards are those created by the execution of one instruction, and seen by the instruction fetch of another instruction. Table 7-2 lists instruction hazards.

Producer	$\rightarrow$	Consumer	Hazard On	"Typical" Spacing (Cycles)
TLBWR, TLBWI	$\rightarrow$	Instruction fetch using new TLB entry	TLB entry	5
MTC0	$\rightarrow$	Instruction fetch seeing the new value (including a change to ERL followed by an instruction fetch from the useg segment)	Status	5
MTC0	$\rightarrow$	Instruction fetch seeing the new value	EntryHi <sub>ASID</sub> WatchHi WatchLo	5

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

Producer $\rightarrow$		Consumer	Hazard On	"Typical" Spacing (Cycles)
Instruction stream writes	$\rightarrow$	Instruction fetch seeing the new instruction stream	Cache entries	Unbounded
CACHE	$\rightarrow$	Instruction fetch seeing the new instruction stream	Cache entries	5

#### **Table 7-2 Instruction Hazards**

# 7.3 Hazard Clearing Instructions

Table 7-3 lists the instructions designed to eliminate hazards.

Mnemonic	Function
EHB	Clear execution hazard
JALR.HB	Clear both execution and instruction hazards
JR.HB	Clear both execution and instruction hazards
SSNOP	Superscalar No Operation
SYNCI	Synchronize caches after instruction stream write

## 7.3.1 Instruction Encoding

The EHB instruction is encoded using a variant of the NOP/SSNOP encoding. This encoding was chosen for compatibility with the Release 1 SSNOP instruction, such that existing software may be modified to be compatible with both Release 1 and Release 2 implementations. See the EHB instruction description for additional information.

The JALR.HB and JR.HB instructions are encoding using bit 10 of the *hint* field of the JALR and JR instructions. These encodings were chosen for compatibility with existing MIPS implementations, including many which pre-date the MIPS32 architecture. Because a pipeline flush clears hazards on most early implementations, the JALR.HB or JR.HB instructions can be included in existing software for backward and forward compatibility. See the JALR.HB and JR.HB instructions for additional information.

The SYNCI instruction is encoded using a new encoding of the REGIMM opcode. This encoding was chosen because it causes a Reserved Instruction exception on all Release 1 implementations. As such, kernel software running on processors that don't implement Release 2 can emulate the function using the CACHE instruction.

# Coprocessor 0 Registers

The Coprocessor 0 (CP0) registers provide the interface between the ISA and the PRA. Each register is discussed below, with the registers presented in numerical order, first by register number, then by select field number.

# 8.1 Coprocessor 0 Register Summary

Table 8-1 lists the CP0 registers in numerical order. The individual registers are described later in this document. If the compliance level is qualified (e.g., "*Required* (TLB MMU)"), it applies only if the qualifying condition is true. The Sel column indicates the value to be used in the field of the same name in the MFC0 and MTC0 instructions.

	Register Number	Sel <sup>1</sup>	Register Name	Function	Reference	Compliance Level
	0	0	Index	Index into the TLB array	Section 8.3 on page 57	Required (TLB MMU); Optional (others)
	1	0	Random	Randomly generated index into the TLB array	Section 8.4 on page 58	Required (TLB MMU); Optional (others)
	2	0	EntryLo0	Low-order portion of the TLB entry for even-numbered virtual pages	Section 8.5 on page 59	Required (TLB MMU); Optional (others)
	3	0	EntryLo1	Low-order portion of the TLB entry for odd-numbered virtual pages	Section 8.5 on page 59	Required (TLB MMU); Optional (others)
	4	0	Context	Pointer to page table entry in memory	Section 8.6 on page 63	Required (TLB MMU); Optional (others)
	4	1	ContextConfig	Context and XContext register configuration	SmartMIPS ASE Specification	Required (SmartMIPS ASE Only)
	5	0	PageMask	Control for variable page size in TLB entries	Section 8.7 on page 64	Required (TLB MMU); Optional (others)
	5	1	PageGrain	Control for small page support	Section 8.8 on page 66 and SmartMIPS ASE Specification	Required (SmartMIPS ASE); Optional (Release 2)
-	6	0	Wired	Controls the number of fixed ("wired") TLB entries	Section 8.9 on page 68	Required (TLB MMU); Optional (others)

 Table 8-1 Coprocessor 0 Registers in Numerical Order

Register Number	Sel <sup>1</sup>	Register Name	Function	Reference	Compliance Level
7	0	HWREna	Enables access via the RDHWR instruction to selected hardware registers	Section 8.10 on page 69	Required (Release 2)
7	1-7		Reserved for future extensions		Reserved
8	0	BadVAddr	Reports the address for the most recent address-related exception	Section 8.11 on page 70	Required
9	0	Count	Processor cycle count	Section 8.12 on page 71	Required
9	6-7		Available for implementation dependent user	Section 8.13 on page 71	Implementation Dependent
10	0	EntryHi	High-order portion of the TLB entry	Section 8.14 on page 72	Required (TLB MMU); Optional (others)
11	0	Compare	Timer interrupt control	Section 8.15 on page 74	Required
11	6-7		Available for implementation dependent user	Section 8.16 on page 74	Implementation Dependent
12	0	Status	Processor status and control	Section 8.17 on page 75	Required
12	1	IntCtl	Interrupt system status and control	Section 8.18 on page 82	Required (Release 2)
12	2	SRSCtl	Shadow register set status and control	Section 8.19 on page 84	Required (Release 2)
12	3	SRSMap	Shadow set IPL mapping	Section 8.20 on page 86	Required (Release 2 and shadow sets implemented)
13	0	Cause	Cause of last general exception	Section 8.21 on page 87	Required
14	0	EPC	Program counter at last exception	Section 8.22 on page 91	Required
15	0	PRId	Processor identification and revision	Section 8.23 on page 92	Required
15	1	EBase	Exception vector base register	Section 8.24 on page 93	Required (Release 2)
16	0	Config	Configuration register	Section 8.25 on page 95	Required
16	1	Config1	Configuration register 1	Section 8.26 on page 97	Required
16	2	Config2	Configuration register 2	Section 8.27 on page 101	Optional
16	3	Config3	Configuration register 3	Section 8.28 on page 104	Optional
16	6-7		Available for implementation dependent user	Section 8.29 on page 106	Implementation Dependent

Table 8-1	Coprocessor (	) Registers in	Numerical	Order
Table 0-1	Coprocessor v	, we gisters m	1 vuinei ieai	oruci

I

Register Number	Sel <sup>1</sup>	Register Name	Function	Reference	Compliance Level	
17	0	LLAddr	Load linked address	Section 8.30 on page 107	Optional	
18	0-n	WatchLo	Watchpoint address	Section 8.31 on page 108	Optional	
19	0-n	WatchHi	Watchpoint control	Section 8.32 on page 110	Optional	
20	0		XContext in 64-bit implementations		Reserved	
21	all		Reserved for future extensions		Reserved	
22	all		Available for implementation dependent use	Section 8.33 on page 112	Implementation Dependent	
23	0	Debug	EJTAG Debug register	EJTAG Specification	Optional	
23	1	TraceControl	PDtrace control register	PDtrace Specification	Optional	
23	2	TraceControl2	PDtrace control register	PDtrace Specification	Optional	
23	3	UserTraceData	PDtrace control register	PDtrace Specification	Optional	
23	4	TraceBPC	PDtrace control register	PDtrace Specification	Optional	
24	0	DEPC	Program counter at last EJTAG debug exception	EJTAG Specification	Optional	
25	0-n	PerfCnt	Performance counter interface	Section 8.36 on page 115	Recommended	
26	0	ErrCtl	Parity/ECC error control and status	Section 8.37 on page 118	Optional	
27	0-3	CacheErr	Cache parity error control and status	Section 8.38 on page 119	Optional	
28	even selects	TagLo	Low-order portion of cache tag interface	Section 8.39 on page 120	Required (Cache)	
28	odd selects	DataLo	Low-order portion of cache data interface	Section 8.40 on page 121	Optional	
29	even selects	TagHi	High-order portion of cache tag interface	Section 8.41 on page 122	Required (Cache)	
29	odd selects	DataHi	High-order portion of cache data interface	Section 8.42 on page 123	Optional	
30	0	ErrorEPC	Program counter at last error	Section 8.43 on page 124	Required	
31	0	DESAVE	EJTAG debug exception save register	EJTAG Specification	Optional	

Table 8-1 Coprocessor 0 Registers in Numeric
----------------------------------------------

-

I

I

1. Any select (Sel) value not explicitly noted as available for implementation-dependent use is reserved for future use by the Architecture.

# 8.2 Notation

For each register described below, field descriptions include the read/write properties of the field, and the reset state of the field. For the read/write properties of the field, the following notation is used:

Read/Write Notation	Hardware Interpretation	Software Interpretation
	A field in which all bits are readable and writabl Hardware updates of this field are visible by soft	e by software and, potentially, by hardware.
R/W	visible by hardware read. If the Reset State of this field is "Undefined", eith before the first read will return a predictable valu definition of <b>UNDEFINED</b> behavior.	her software or hardware must initialize the value i. This should not be confused with the formal
R	A field which is either static or is updated only by hardware. If the Reset State of this field is either "0", "Preset", or "Externally Set", hardware initializes this field to zero or to the appropriate state, respectively, on powerup. The term "Preset" is used to suggest that the processor establishes the appropriate state, whereas the term "Externally Set" is used to suggest that the state is established via an external source (e.g., personality pins or initialization bit stream). These terms are suggestions only, and are not intended to act as a requirement on the implementation. If the Reset State of this field is "Undefined", hardware updates this field only under those conditions specified in the description of the field.	A field to which the value written by software is ignored by hardware. Software may write any value to this field without affecting hardware behavior. Software reads of this field return the last value updated by hardware. If the Reset State of this field is "Undefined", software reads of this field result in an <b>UNPREDICTABLE</b> value except after a hardware update done under the conditions specified in the description of the field.
0	A field which hardware does not update, and for which hardware can assume a zero value.	A field to which the value written by software must be zero. Software writes of non-zero values to this field may result in <b>UNDEFINED</b> behavior of the hardware. Software reads of this field return zero as long as all previous software writes are zero. If the Reset State of this field is "Undefined", software must write this field with zero before it is guaranteed to read as zero.

## Table 8-2 Read/Write Bit Field Notation

## 8.3 Index Register (CP0 Register 0, Select 0)

Compliance Level: Required for TLB-based MMUs; Optional otherwise.

The *Index* register is a 32-bit read/write register which contains the index used to access the TLB for TLBP, TLBR, and TLBWI instructions. The width of the index field is implementation-dependent as a function of the number of TLB entries that are implemented. The minimum value for TLB-based MMUs is Ceiling(Log2(TLBEntries)). For example, six bits are required for a TLB with 48 entries).

The operation of the processor is **UNDEFINED** if a value greater than or equal to the number of TLB entries is written to the *Index* register.

Figure 8-1 shows the format of the Index register; Table 8-3 describes the Index register fields.

## **Figure 8-1 Index Register Format**

31	n	n-1 0
Р	0	Index

Fields				Read/			
Name	Bits	Description		Write	Reset State	Compliance	
		Probe Failu execution of a TLB mate	re. Hardware writes this bit during f the TLBP instruction to indicate whether ch occurred:				
D	21	Encoding	Meaning		Undefined	D a suciona d	
Ρ	51	0	A match occurred, and the Index field contains the index of the matching entry	ĸ	Undefined	Required	
		1	No match occurred and the Index field is <b>UNPREDICTABLE</b>				
0	30n	Must be wi	itten as zero; returns zero on read.	0	0	Reserved	
Index	n-10	TLB index. Software writes this field to provide the index to the TLB entry referenced by the TLBR and TLBWI instructions. Hardware writes this field with the index of the matching TLB entry during execution of the TLBP instruction. If the TLBP fails to find a match, the contents of this field are <b>UNPREDICTABLE</b> .		R/W	Undefined	Required	

## **Table 8-3 Index Register Field Descriptions**

# 8.4 Random Register (CP0 Register 1, Select 0)

Compliance Level: Required for TLB-based MMUs; Optional otherwise.

The *Random* register is a read-only register whose value is used to index the TLB during a TLBWR instruction. The width of the Random field is calculated in the same manner as that described for the *Index* register above.

The value of the register varies between an upper and lower bound as follow:

- A lower bound is set by the number of TLB entries reserved for exclusive use by the operating system (the contents of the *Wired* register). The entry indexed by the *Wired* register is the first entry available to be written by a TLB Write Random operation.
- An upper bound is set by the total number of TLB entries minus 1.

Within the required constraints of the upper and lower bounds, the manner in which the processor selects values for the Random register is implementation-dependent.

The processor initializes the *Random* register to the upper bound on a Reset Exception, and when the *Wired* register is written.

Figure 8-2 shows the format of the Random register; Table 8-4 describes the Random register fields.

## Figure 8-2 Random Register Format

31	n	n-1 0
	0	Random

Table 8-4 Random Register Fi	eld Descriptions
------------------------------	------------------

Fiel	lds		Read/		
Name	Bits	Description	Write	Reset State	Compliance
0	31n	Must be written as zero; returns zero on read.	0	0	Reserved
Random	n-10	TLB Random Index	R	TLB Entries - 1	Required

# 8.5 EntryLo0, EntryLo1 (CP0 Registers 2 and 3, Select 0)

Compliance Level: EntryLo0 is Required for a TLB-based MMU; Optional otherwise.

Compliance Level: EntryLo1 is Required for a TLB-based MMU; Optional otherwise.

The pair of *EntryLo* registers act as the interface between the TLB and the TLBP, TLBR, TLBWI, and TLBWR instructions. *EntryLo0* holds the entries for even pages and *EntryLo1* holds the entries for odd pages.

Software may determine the value of PABITS by writing all ones to the *EntryLo0* or *EntryLo1* registers and reading the value back. Bits read as "1" from the PFN field allow software to determine the boundary between the PFN and Fill fields to calculate the value of PABITS.

The contents of the *EntryLo0* and *EntryLo1* registers are not defined after an address error exception and some fields may be modified by hardware during the address error exception sequence. Software writes of the *EntryHi* register (via MTC0) do not cause the implicit update of address-related fields in the *BadVAddr* or *Context* registers.

For Release 1 of the Architecture, Figure 8-3 shows the format of the *EntryLo0* and *EntryLo1* registers; Table 8-5 describes the *EntryLo0* and *EntryLo1* register fields. For Release 2 of the Architecture, Figure 8-4 shows the format of the *EntryLo0* and *EntryLo1* registers; Table 8-6 describes the *EntryLo0* and *EntryLo1* register fields.

#### Figure 8-3 EntryLo0, EntryLo1 Register Format in Release 1 of the Architecture

31 30	29 6	5	3	2	1	0
Fill	PFN		С	D	V	G

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
Fill	3130	These bits are ignored on write and return zero on read. The boundaries of this field change as a function of the value of <i>PABITS</i> . See Table 8-7 for more information.	R	0	Required
PFN	296	Page Frame Number. Corresponds to bits <i>PABITS</i> -112 of the physical address, where <i>PABITS</i> is the width of the physical address in bits. The boundaries of this field change as a function of the value of <i>PABITS</i> . See Table 8-7 for more information.	R/W	Undefined	Required
С	53	Coherency attribute of the page. See Table 8-8 below.	R/W	Undefined	Required
D	2	"Dirty" bit, indicating that the page is writable. If this bit is a one, stores to the page are permitted. If this bit is a zero, stores to the page cause a TLB Modified exception. Kernel software may use this bit to implement paging algorithms that require knowing which pages have been written. If this bit is always zero when a page is initially mapped, the TLB Modified exception that results on any store to the page can be used to update kernel data structures that indicate that the page was actually written.	R/W	Undefined	Required
v	1	Valid bit, indicating that the TLB entry, and thus the virtual page mapping are valid. If this bit is a one, accesses to the page are permitted. If this bit is a zero, accesses to the page cause a TLB Invalid exception.	R/W	Undefined	Required

## Table 8-5 EntryLo0, EntryLo1 Register Field Descriptions in Release 1 of the Architecture

## Table 8-5 EntryLo0, EntryLo1 Register Field Descriptions in Release 1 of the Architecture

Fields			Read/			
Name	Bits	Description	Write	Reset State	Compliance	
G	0	Global bit. On a TLB write, the logical AND of the G bits from both EntryLo0 and EntryLo1 becomes the G bit in the TLB entry. If the TLB entry G bit is a one, ASID comparisons are ignored during TLB matches. On a read from a TLB entry, the G bits of both EntryLo0 and EntryLo1 reflect the state of the TLB G bit.	R/W	Undefined	Required (TLB MMU)	

## Figure 8-4 EntryLo0, EntryLo1 Register Format in Release 2 of the Architecture

31 30	29 6	5	3	2	1	0
Fill	PFN	(	2	D	V	G

## Table 8-6 EntryLo0, EntryLo1 Register Field Descriptions in Release 2 of the Architecture

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
Fill	3130	These bits are ignored on write and return zero on read. The boundaries of this field change as a function of the value of PABITS. See Table 8-7 for more information.	R	0	Required
PFN	296	Page Frame Number. This field contains the physical page number corresponding to the virtual page. If the processor is enabled to support 1KB pages (Config $_{SP}$ = 1 and PageGrain <sub>ESP</sub> = 1), the PFN field corresponds to bits 3310 of the physical address (the field is shifted left by 2 bits relative to the Release 1 definition to make room for PA <sub>1110</sub> ). If the processor is not enabled to support 1KB pages (Config $_{SP}$ = 0 or PageGrain <sub>ESP</sub> = 0), the PFN field corresponds to bits 3512 of the physical address. The boundaries of this field change as a function of the value of PABITS. See Table 8-7 for more information.	R/W	Undefined	Required
С	53	The definition of this field is unchanged from Release 1. See Table 8-5 above and Table 8-8 below.	R/W	Undefined	Required
D	2	The definition of this field is unchanged from Release 1. See Table 8-5 above.	R/W	Undefined	Required
V	1	The definition of this field is unchanged from Release 1. See Table 8-5 above.	R/W	Undefined	Required
G	0	The definition of this field is unchanged from Release 1. See Table 8-5 above.	R/W	Undefined	Required (TLB MMU)

Table 8-7 shows the movement of the Fill and PFN fields as a function of 1KB page support enabled, and the value of *PABITS*. Note that in implementations of Release 1 of the Architecture, there is no support for 1KB pages, so only the first row of the table applies to Release 1.

1KB Page Support		Corresponding Ra	Release 2	
Enabled?	PABITS Value	Fill Field	PFN Field	Required?
No	36 ≥ PABITS > 12	31(30-(36- <i>PABITS</i> ) ) Example: 3130 if <i>PABITS</i> = 36 317 if <i>PABITS</i> = 13	(29-(36- <i>PABITS</i> ))6 Example: 296 if <i>PABITS</i> = 36 66 if <i>PABITS</i> = 13 EntryLo <sub>296</sub> = PA <sub>3512</sub>	No
Yes	34 ≥ PABITS > 10	31(30-(34- <i>PABITS</i> ) ) Example: 3130 if <i>PABITS</i> = 34 317 if <i>PABITS</i> = 11	(29-(34- <i>PABITS</i> ))6 Example: 296 if <i>PABITS</i> = 34 66 if <i>PABITS</i> = 11 EntryLo <sub>296</sub> = PA <sub>3310</sub>	Yes

## Table 8-7 EntryLo Field Widths as a Function of PABITS

## **Programming Note:**

In implementations of Release 2 of the Architecture, the PFN field of both the *EntryLo0* and *EntryLo1* registers must be written with zero and the TLB must be flushed before each instance in which the value of the *PageGrain* register is changed. This operation must be carried out while running in an unmapped address space. The operation of the processor is **UNDEFINED** if this sequence is not done.

Table 8-8 lists the encoding of the C field of the *EntryLo0* and *EntryLo1* registers and the K0 field of the *Config* register. An implementation may choose to implement a subset of the cache coherency attributes shown, but must implement at least encodings 2 and 3 such that software can always depend on these encodings working appropriately. In other cases, the operation of the processor is **UNDEFINED** if software specifies an unimplemented encoding.

 Table 8-8 lists the required and optional encodings for the coherency attributes.

#### Table 8-8 Cache Coherency Attributes

C(5:3) Value	Cache Coherency Attributes With Historical Usage	Compliance
0	Available for implementation dependent use	Optional
1	Available for implementation dependent use	Optional
2	Uncached	Required
3	Cacheable	Required
4	Available for implementation dependent use	Optional

# Table 8-8 Cache Coherency Attributes

C(5:3) Value	Cache Coherency Attributes With Historical Usage	Compliance
5	Available for implementation dependent use	Optional
6	Available for implementation dependent use	Optional
7	Available for implementation dependent use	Optional

## 8.6 Context Register (CP0 Register 4, Select 0)

Compliance Level: Required for TLB-based MMUs; Optional otherwise.

The *Context* register is a read/write register containing a pointer to an entry in the page table entry (PTE) array. This array is an operating system data structure that stores virtual-to-physical translations. During a TLB miss, the operating system loads the TLB with the missing translation from the PTE array. The *Context* register duplicates some of the information provided in the *BadVAddr* register, but is organized in such a way that the operating system can directly reference a 16-byte structure in memory that describes the mapping.

A TLB exception (TLB Refill, TLB Invalid, or TLB Modified) causes bits  $VA_{31..13}$  of the virtual address to be written into the *BadVPN2* field of the *Context* register. The *PTEBase* field is written and used by the operating system.

The BadVPN2 field of the *Context* register is not defined after an address error exception and this field may be modified by hardware during the address error exception sequence.

Figure 8-5 shows the format of the Context Register; Table 8-9 describes the Context register fields.

## Figure 8-5 Context Register Format

31	23	22 4	3	0
	PTEBase	BadVPN2		0

Fie	lds		Read/		
Name	Bits	Description	Write	Reset State	Compliance
PTEBase	3123	This field is for use by the operating system and is normally written with a value that allows the operating system to use the <i>Context</i> Register as a pointer into the current PTE array in memory.	R/W	Undefined	Required
BadVPN2	224	This field is written by hardware on a TLB exception. It contains bits $VA_{3113}$ of the virtual address that caused the exception.	R	Undefined	Required
0	30	Must be written as zero; returns zero on read.	0	0	Reserved

## Table 8-9 Context Register Field Descriptions

# 8.7 PageMask Register (CP0 Register 5, Select 0)

Compliance Level: Required for TLB-based MMUs; Optional otherwise.

The *PageMask* register is a read/write register used for reading from and writing to the TLB. It holds a comparison mask that sets the variable page size for each TLB entry, as shown in Table 8-11. Figure 8-6 shows the format of the *PageMask* register; Table 8-10 describes the *PageMask* register fields.

## Figure 8-6 PageMask Register Format

31 29	28 1	3 12 11		0
0	Mask	MaskX	0	

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
Mask	2813	The Mask field is a bit mask in which a "1" bit indicates that the corresponding bit of the virtual address should not participate in the TLB match.	R/W	Undefined	Required
MaskX	1211	In Release 2 of the Architecture, the MaskX field is an extension to the Mask field to support 1KB pages with definition and action analogous to that of the Mask field, defined above. If 1KB pages are enabled (Config $3_{SP} = 1$ and PageGrain <sub>ESP</sub> = 1), these bits are writable and readable, and their values are copied to and from the TLB entry on a TLB write or read, respectivly. If 1KB pages are not enabled (Config $3_{SP} = 0$ or PageGrain <sub>ESP</sub> = 0), these bits are not writable, return zero on read, and the effect on the TLB entry on a write is as if they were written with the value 2#11. In Release 1 of the Architecture, these bits must be written as zero, return zero on read, and have no effect on the virtual address translation.	R/W	0 (See Description)	Required (Release 2)
0	3129, 100	Ignored on write; returns zero on read.	R	0	Required

#### **Table 8-10 PageMask Register Field Descriptions**

# Table 8-11 Values for the Mask and MaskX<sup>1</sup> Fields of the PageMask Register

		Bit																
Page Size	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12 <sup>1</sup>	11 <sup>1</sup>
1 KByte	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4 KBytes	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1
16 KBytes	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1
64 KBytes	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1
256 KBytes	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
1 MByte	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1
4 MByte	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1
		Bit																
-----------	----	-----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	-----------------	-----------------
Page Size	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12 <sup>1</sup>	11 <sup>1</sup>
16 MByte	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
64 MByte	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
256 MByte	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

## Table 8-11 Values for the Mask and MaskX<sup>1</sup> Fields of the PageMask Register

1. PageMask<sub>12..11</sub> = PaskMask<sub>MaskX</sub> exists only on implementations of Release 2 of the architecture and are treated as if they had the value 2#11 if 1K pages are not enabled (Config3<sub>SP</sub> = 0 or PageGrain<sub>ESP</sub> = 0).

It is implementation dependent how many of the encodings described in Table 8-11 are implemented. All processors must implement the 4KB page size. If a particular page size encoding is not implemented by a processor, a read of the *PageMask* register must return zeros in all bits that correspond to encodings that are not implemented, thereby potentially returning a value different than that written by software.

Software may determine which page sizes are supported by writing all ones to the *PageMask* register, then reading the value back. If a pair of bits reads back as ones, the processor implements that page size. The operation of the processor is **UNDEFINED** if software loads the Mask field with a value other than one of those listed in Table 8-11, even if the hardware returns a different value on read. Hardware may depend on this requirement in implementing hardware structures

#### **Programming Note:**

In implementations of Release 2 of the Architecture, the MaskX field of the *PageMask* register must be written with 2#11 and the TLB must be flushed before each instance in which the value of the *PageGrain* register is changed. This operation must be carried out while running in an unmapped address space. The operation of the processor is **UNDE-FINED** if this sequence is not done.

# 8.8 PageGrain Register (CP0 Register 5, Select 1)

**Compliance Level:** *Required* for implementations of Release 2 of the Architecture that include TLB-based MMUs and support 1KB pages; *Optional* otherwise.

The *PageGrain* register is a read/write register used for enabling 1KB page support. The *PageGrain* register is present in both the SmartMIPS<sup>TM</sup> ASE, and in Release 2 of the Architecture, although there are no bits in common between the two uses of this register. As such, the description below only describes the fields relevant to Release 2 of the Architecture. In implementations of both Release 2 of the Architecture and the SmartMIPS<sup>TM</sup> ASE, the ASE definitions take precedence and none of the Release 2 fields described below are present. Figure 8-7 shows the format of the *PageMask* register; Table 8-12 describes the *PageMask* register fields.

### **Figure 8-7 PageGrain Register Format**

31 30	29 28	27 13	12 8	7 0
ASE	ELPA ESP	0	ASE	0

Fields				Dood/		
Name	Bits		Description	Write	Reset State	Compliance
ASE	3130,1 28	These fields ASE and are of the Archite implements t	are control features of the SmartMIPS <sup>TM</sup> not used in implementations of Release 2 ecture unless such an implementation also he SmartMIPS <sup>TM</sup> ASE.	0	0	Required
ELPA	29	Used to enab MIPS64 prod This bit is ig	le support for large physical addresses in cessors; not used by MIPS32 processors. nored on write and returns zero on read.	R	0	Required
		Enables supp	oort for 1KB pages.			
		Encoding	Meaning			
		0	1KB page support is not enabled			
		1	1KB page support is enabled			
ESP	28	<ul> <li>If this bit is a coprocessor ' <ul> <li>The PFN f</li> <li>registers h</li> <li>(the field i</li> <li>definition)</li> </ul> </li> <li>The Mask writable and field to for entry.</li> <li>The VPN2 and bits 12</li> <li>The virtua to reflect t</li> <li>If Config3sp this bit is ign</li> </ul>	a 1, the following changes occur to D registers: held of the <i>EntryLo0</i> and <i>EntryLo1</i> olds the physical address down to bit 10 s shifted left by 2 bits from the Release 1 X field of the <i>PageMask</i> register is hd is concatenated to the right of the Mask m the "don't care" mask for the TLB EX field of the <i>EntryHi</i> register is writable 211 of the virtual address. I address translation algorithm is modified he smaller page size. = 0, 1KB pages are not implemented, and ored on write and returns zero on read.	R/W	0	Required
0	2713,	Must be writ	ten as zero: returns zero on read	0	0	Reserved
v	70		ten us zero, returns zero on retur.			itesei veu

### Table 8-12 PageGrain Register Field Descriptions

#### **Programming Note:**

In implementations of Release 2 of the Architecture, the following fields must be written with the specified values, and the TLB must be flushed before each instance in which the value of the PageGrain register is changed. This operation must be carried out while running in an unmapped address space. The operation of the processor is **UNDE-FINED** if this sequence is not done.

Field	<b>Required Value</b>
EntryLo0 <sub>PFN</sub> , EntryLo1 <sub>PFN</sub>	0
EntryLo0 <sub>PFNX</sub> , EntryLo1 <sub>PFNX</sub>	0
PageMask <sub>MaskX</sub>	2#11
EntryHi <sub>VPN2X</sub>	0

Note also that if PageGrain is changed, a hazard may be created between the instruction that writes PageGrain and a subsequent CACHE instruction. This hazard must be cleared using the EHB instruction.

### 8.9 Wired Register (CP0 Register 6, Select 0)

Compliance Level: Required for TLB-based MMUs; Optional otherwise.

The *Wired* register is a read/write register that specifies the boundary between the wired and random entries in the TLB as shown in Figure 8-8.



The width of the Wired field is calculated in the same manner as that described for the *Index* register. Wired entries are fixed, non-replaceable entries which are not overwritten by a TLBWR instruction. Wired entries can be overwritten by a TLBWI instruction.

The *Wired* register is set to zero by a Reset Exception. Writing the *Wired* register causes the *Random* register to reset to its upper bound.

The operation of the processor is **UNDEFINED** if a value greater than or equal to the number of TLB entries is written to the *Wired* register.

Figure 8-8 shows the format of the Wired register; Table 8-13 describes the Wired register fields.

### **Figure 8-9 Wired Register Format**

31	n	n-1 0
	0	Wired

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
0	31n	Must be written as zero; returns zero on read.	0	0	Reserved
Wired	n-10	TLB wired boundary	R/W	0	Required

#### **Table 8-13 Wired Register Field Descriptions**

### 8.10 HWREna Register (CP0 Register 7, Select 0)

Compliance Level: Required (Release 2).

The *HWREna* register contains a bit mask that determines which hardware registers are accessible via the RDHWR instruction.

Figure 8-10 shows the format of the *HWREna* Register; Table 8-14 describes the *HWREna* register fields.

#### Figure 8-10 HWREna Register Format

31	4	3	0
0 0000 0000 0000 0000 0000 0000 0000			Mask

#### Table 8-14 HWREna Register Field Descriptions

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
0	314	Must be written with zero; returns zero on read	0	0	Reserved
Mask	30	Each bit in this field enables access by the RDHWR instruction to a particular hardware register (which may not be an actual register). If bit 'n' in this field is a 1, access is enabled to hardware register 'n'. If bit 'n' of this field is a 0, access is disabled. See the RDHWR instruction for a list of valid hardware registers.	R/W	0	Required

Privileged software may determine which of the hardware registers are accessible by the RDHWR instruction. In doing so, a register may be virtualized at the cost of handling a Reserved Instruction Exception, interpreting the instruction, and returning the virtualized value. For example, if it is not desirable to provide direct access to the *Count* register, access to that register may be individually disabled and the return value can be virtualized by the operating system.

## 8.11 BadVAddr Register (CP0 Register 8, Select 0)

### Compliance Level: Required.

The *BadVAddr* register is a read-only register that captures the most recent virtual address that caused one of the following exceptions:

- Address error (AdEL or AdES)
- TLB Refill
- TLB Invalid (TLBL, TLBS)
- TLB Modified

The *BadVAddr* register does not capture address information for cache or bus errors, or for Watch exceptions, since none is an addressing error.

Figure 8-11 shows the format of the BadVAddr register; Table 8-15 describes the BadVAddr register fields.

### Figure 8-11 BadVAddr Register Format

31	0
BadVAddr	

### Table 8-15 BadVAddr Register Field Descriptions

Fields			Read/			
Name	Bits	Description	Write	Reset State	Compliance	
BadVAddr	310	Bad virtual address	R	Undefined	Required	

## 8.12 Count Register (CP0 Register 9, Select 0)

### Compliance Level: Required.

The Count register acts as a timer, incrementing at a constant rate, whether or not an instruction is executed, retired, or any forward progress is made through the pipeline. The rate at which the counter increments is implementation dependent, and is a function of the pipeline clock of the processor, not the issue width of the processor.

The Count register can be written for functional or diagnostic purposes, including at reset or to synchronize processors.

Figure 8-12 shows the format of the Count register; Table 8-16 describes the Count register fields.

### **Figure 8-12 Count Register Format**

31	0
	Count

### **Table 8-16 Count Register Field Descriptions**

Fields			Read/			
Name	Bits	Description	Write	Reset State	Compliance	
Count	310	Interval counter	R/W	Undefined	Required	

### 8.13 Reserved for Implementations (CP0 Register 9, Selects 6 and 7)

Compliance Level: Optional: Implementation Dependent.

CP0 register 9, Selects 6 and 7 are reserved for implementation dependent use and are not defined by the architecture.

## 8.14 EntryHi Register (CP0 Register 10, Select 0)

Compliance Level: Required for TLB-based MMU; Optional otherwise.

The EntryHi register contains the virtual address match information used for TLB read, write, and access operations.

A TLB exception (TLB Refill, TLB Invalid, or TLB Modified) causes bits  $VA_{31..13}$  of the virtual address to be written into the VPN2 field of the *EntryHi* register. An implementation of Release 2 of the Architecture which supports 1KB pages also writes  $VA_{12..11}$  into the VPN2X field of the EntryHi register. A TLBR instruction writes the *EntryHi* register with the corresponding fields from the selected TLB entry. The ASID field is written by software with the current address space identifier value and is used during the TLB comparison process to determine TLB match.

Because the ASID field is overwritten by a TLBR instruction, software must save and restore the value of ASID around use of the TLBR. This is especially important in TLB Invalid and TLB Modified exceptions, and in other memory management software.

The VPNX2 and VPN2 fields of the *EntryHi* register are not defined after an address error exception and these fields may be modified by hardware during the address error exception sequence.Software writes of the EntryHi register (via MTC0) do not cause the implicit write of address-related fields in the *BadVAddr,Context* registers.

Figure 8-13 shows the format of the EntryHi register; Table 8-17 describes the EntryHi register fields.

### Figure 8-13 EntryHi Register Format

31 13	12 11	10 8	7 0
VPN2	VPN2X	0	ASID

Fields			Read/		
Name	Bits	Description	Write	<b>Reset State</b>	Compliance
VPN2	3113	$VA_{31,.13}$ of the virtual address (virtual page number / 2). This field is written by hardware on a TLB exception or on a TLB read, and is written by software before a TLB write.	R/W	Undefined	Required
VPN2X	1211	In Release 2 of the Architecture, the VPN2X field is an extension to the VPN2 field to support 1KB pages. These bits are not writable by either hardware or software unless $Config3_{SP} = 1$ and $PageGrain_{ESP} = 1$ . If enabled for write, this field contains $VA_{1211}$ of the virtual address and is written by hardware on a TLB exception or on a TLB read, and is by software before a TLB write. If writes are not enabled, and in implementations of Release 1 of the Architecture, this field must be written with zero and returns zeros on read.	R/W	0	Required (Release 2 and 1KB Page Support)
0	108	Must be written as zero; returns zero on read.	0	0	Reserved
ASID	70	Address space identifier. This field is written by hardware on a TLB read and by software to establish the current ASID value for TLB write and against which TLB references match each entry's TLB ASID field.	R/W	Undefined	Required (TLB MMU)

### Table 8-17 EntryHi Register Field Descriptions

### **Programming Note:**

In implementations of Release 2 of the Architecture, the VPN2X field of the *EntryHi* register must be written with zero and the TLB must be flushed before each instance in which the value of the *PageGrain* register is changed. This

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

operation must be carried out while running in an unmapped address space. The operation of the processor is **UNDE-FINED** if this sequence is not done.

# 8.15 Compare Register (CP0 Register 11, Select 0)

### Compliance Level: Required.

The *Compare* register acts in conjunction with the *Count* register to implement a timer and timer interrupt function. The *Compare* register maintains a stable value and does not change on its own.

When the value of the *Count* register equals the value of the *Compare* register, an interrupt request is combined in an implementation-dependent way with hardware interrupt 5 to set interrupt bit IP(7) in the *Cause* register. This causes an interrupt as soon as the interrupt is enabled.

For diagnostic purposes, the *Compare* register is a read/write register. In normal use however, the *Compare* register is write-only. Writing a value to the *Compare* register, as a side effect, clears the timer interrupt. Figure 8-14 shows the format of the *Compare* register; Table 8-18 describes the Compare register fields.

### **Figure 8-14 Compare Register Format**

31	0
	Compare

### **Table 8-18 Compare Register Field Descriptions**

Fiel	Fields				
Name	Bits	Description	Write	Reset State	Compliance
Compare	310	Interval count compare value	R/W	Undefined	Required

## 8.16 Reserved for Implementations (CP0 Register 11, Selects 6 and 7)

### Compliance Level: Optional: Implementation Dependent.

CP0 register 11, Selects 6 and 7 are reserved for implementation dependent use and are not defined by the architecture.

### 8.17 Status Register (CP Register 12, Select 0)

### Compliance Level: Required.

The *Status* register is a read/write register that contains the operating mode, interrupt enabling, and the diagnostic states of the processor. Fields of this register combine to create operating modes for the processor. Refer to Chapter 3, "MIPS32 Operating Modes," on page 9 for a discussion of operating modes, and Section Section 5.1, "Interrupts" on page 23 for a discussion of interrupt modes.

Figure 8-15 shows the format of the Status register; Table 8-19 describes the Status register fields.

### **Figure 8-15 Status Register Format**

	31	28	27	26	25	24	23	22	21	20	19	18	17 16	15		10	9	8	7	6	5	4	3	2	1	0
	CU3	.CU0	RP	FR	RE	MX	PX	BEV	TS	SR	NMI	0	Impl		IM7IM2		IM1.	.IM0	KΧ	SX	UX	UM	R0	ERL	EXL	IE
_															IPL							KS	U			

Fields				Read/		
Name	Bits		Description	Write	Reset State	Compliance
		Controls acc respectively:	ess to coprocessors 3, 2, 1, and 0,			
		Encoding	Meaning			
		0	Access not allowed		Undefined	
		1	Access allowed			
CU (CU3 CU0)	3128	Coprocessor running in K the state of th In Release 2 implementat execution of those encode the CU1 enal future use by If there is no correspondin as zero.	0 is always usable when the processor is ernel Mode or Debug Mode, independent of he $CU_0$ bit. of the Architecture, and for 64-bit ions of Release 1 of the Architecture, all floating point instructions, including d with the COP1X opcode, is controlled by ble. CU3 is no longer used and is reserved for the Architecture. provision for connecting a coprocessor, the g CU bit must be ignored on write and read	R/W		Required for all implemented coprocessors
RP	27	Enables redu The specific dependent. If this bit is r and read as z must be zero performance	ced power mode on some implementations. operation of this bit is implementation not implemented, it must be ignored on write ero. If this bit is implemented, the reset state so that the processor starts at full	R/W	0	Optional

### Table 8-19 Status Register Field Descriptions

<b>Table 8-19 Status Register</b>	Field Descriptions
-----------------------------------	--------------------

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
		In Release 1 of the Architecture, only MIPS64 processors could implement a 64-bit floating point unit. In Release 2 of the Architecture, both MIPS32 and MIPS64 processors can implement a 64-bit floating point unit. This bit is used to control the floating point register mode for 64-bit floating point units:			
		Encoding Meaning			
		Floating point registers can contain any         0       32-bit datatype. 64-bit datatypes are stored         in even-odd pairs of registers.			
		1 Floating point registers can contain any datatype		0	
FR	26	This bit must be ignored on write and read as zero under the following conditions:	R		Required
		• No floating point unit is implemented			
		• In a MIPS32 implementation of Release 1 of the Architecture			
		• In an implementation of Release 2 of the Architecture in which a 64-bit floating point unit is not implemented			
		Certain combinations of the FR bit and other state or operations can cause <b>UNPREDICTABLE</b> behavior. See Section Section 3.5.2, "64-bit FPR Enable" on page 10 for a discussion of these combinations.			
		Used to enable reverse-endian memory references while the processor is running in user mode:			
		Encoding Meaning			
		0 User mode uses configured endianness			
RE	25	1 User mode uses reversed endianness	R/W	Undefined	Optional
		Neither Debug Mode nor Kernel Mode nor Supervisor Mode references are affected by the state of this bit.			
		If this bit is not implemented, it must be ignored on write and read as zero.			
MX	24	Enables access to MDMX <sup>TM</sup> resources on MIPS64 processors. Not used by MIPS32 processors. This bit must be ignored on write and read as zero.	R	0	Optional
РХ	23	Enables access to 64-bit operations on MIPS64 processors. Not used by MIPS32 processors. This bit must be ignored on write and read as zero.	R	0	Required

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
BEV	22	Encoding       Meaning         0       Normal         1       Bootstrap    See Section Section 5.2.1, "Exception Vector Locations" on page 32 for details.	R/W	1	Required
TS	21	Indicates that the TLB has detected a match on multiple entries. It is implementation dependent whether this detection occurs at all, on a write to the TLB, or an access to the TLB. In Release 2 of the Architecture, multiple TLB matches may only be reported on a TLB write. When such a detection occurs, the processor initiates a machine check exception and sets this bit. It is implementation dependent whether this condition can be corrected by software. If the condition can be corrected, this bit should be cleared by software before resuming normal operation. See Section 4.9.3 on page 17 for a discussion of software TLB initialization used to avoid a machine check exceeption during processor initialization. If this bit is not implemented, it must be ignored on write and read as zero. Software should not write a 1 to this bit when its value is a 0, thereby causing a 0-to-1 transition. If such a transition is caused by software, it is <b>UNPREDICTABLE</b> whether hardware ignores the write, accepts the write with no side effects, or accepts the write and initiates a machine check exception.	R/W	0	Required if TLB Shutdown is implemented
SR	20	Indicates that the entry through the reset exception vector was due to a Soft Reset:         Encoding       Meaning         0       Not Soft Reset (NMI or Reset)         1       Soft Reset         If this bit is not implemented, it must be ignored on write and read as zero.         Software should not write a 1 to this bit when its value is a 0, thereby causing a 0-to-1 transition. If such a transition is caused by software, it is UNPREDICTABLE whether hardware ignores or accepts the write.	R/W	1 for Soft Reset; 0 otherwise	Required if Soft Reset is implemented

Fields				Read/			
Name	Bits		Description	Write	Reset State	Compliance	
		Indicates tha was due to a	t the entry through the reset exception vector n NMI exception:				
		Encoding	Meaning				
		0	Not NMI (Soft Reset or Reset)				
		1	NMI		1 for NMI 0	Required if	
NMI	19	If this bit is and read as z Software sho a 0, thereby c is caused by hardware igr	not implemented, it must be ignored on write iero. buld not write a 1 to this bit when its value is causing a 0-to-1 transition. If such a transition software, it is <b>UNPREDICTABLE</b> whether lores or accepts the write.	R/W	otherwise	implemented	
0	18	Must be writ	ten as zero; returns zero on read.	0	0	Reserved	
Impl	1716	These bits ar defined by th they must be	e implementation dependent and are not le architecture. If they are not implemented, ignored on write and read as zero.		Undefined	Optional	
		Interrupt Ma hardware int "Interrupts" enabled inter	sk: Controls the enabling of each of the errupts. Refer to Section Section 5.1, on page 23 for a complete discussion of rupts.		Undefined		
		Encoding	Meaning				
IM7IM2	1510	0	Interrupt request disabled	R/W		Required	
		1	Interrupt request enabled				
		In implemen which EIC in these bits tak as the IPL fie	tations of Release 2 of the Architecture in nterrupt mode is enabled (Config $3_{VEIC} = 1$ ), te on a different meaning and are interpreted eld, described below.				
		Interrupt Prie	prity Level.				
IPL	1510	In implemen which EIC in this field is th An interrupt higher than t	tations of Release 2 of the Architecture in hterrupt mode is enabled (Config $3_{VEIC} = 1$ ), he encoded (063) value of the current IPL. will be signaled only if the requested IPL is his value.	R/W	Undefined	Optional (Release 2 and EIC interrupt mode only)	
		If EIC interr these bits tak as the IM7I	upt mode is not enabled (Config $3_{VEIC} = 0$ ), the on a different meaning and are interpreted M2 bits, described above.				

I

Fields				Read/			
Name	Bits		Description	Write	Reset State	Compliance	
		Interrupt Ma software inte "Interrupts" of enabled inter	sk: Controls the enabling of each of the rrupts. Refer to Section Section 5.1, on page 23 for a complete discussion of rupts.				
		Encoding	Encoding Meaning				
IM1IM0	98	0	Interrupt request disabled	R/W	Undefined	Required	
		1	Interrupt request enabled				
		In implemen which EIC ir these bits are system.	tations of Release 2 of the Architecture in hterrupt mode is enabled (Config $3_{VEIC} = 1$ ) writable, but have no effect on the interrup	), pt			
KX	7	Enables acce MIPS proces bit must be i	ess to 64-bit kernel address space on 64-bit sors. Not used by MIPS32 processors. This gnored on write and read as zero.	s R	0	Reserved	
SX	6	Enables acce 64-bit MIPS This bit mus	ess to 64-bit supervisor address space on processors. Not used by MIPS32 processors t be ignored on write and read as zero.	s. R	0	Reserved	
UX	5	Enables acce MIPS proces bit must be i	ess to 64-bit user address space on 64-bit sors Not used by MIPS32 processors. This gnored on write and read as zero.	R	0	Reserved	
		If Supervisor field denotes See Chapter a full discuss field is:	Mode is implemented, the encoding of this the base operating mode of the processor. 3, "MIPS32 Operating Modes," on page 9 for ion of operating modes. The encoding of this	or is			
		Encoding	Meaning			Required if	
KSU	43	2#00	Base mode is Supervisor Mode	R/W	Undefined	Mode is	
		2#01	Base mode is User Mode	10/11		implemented; Optional	
		2#11	Reserved. The operation of the processor is <b>UNDEFINED</b> if this value is written to the KSU field			otherwise	
		Note: This field below.	eld overlaps the UM and R0 fields, describe	d			

I

Fields			Dood/		
Name	Bits	Description	Write	Reset State	Compliance
UM	4	If Supervisor Mode is not implemented, this bit denotes the base operating mode of the processor. See Chapter 3, "MIPS32 Operating Modes," on page 9 for a full discussion of operating modes. The encoding of this bit is:         Encoding       Meaning	R/W	Undefined	Required
		0 Base mode is Kernel Mode			
		I Base mode is User Mode			
		Note: This bit overlaps the KSU field, described above.			
R0	3	If Supervisor Mode is not implemented, this bit is reserved. This bit must be ignored on write and read as zero.	R	0	Reserved
		Note: This bit overlaps the KSU field, described above.			
		Error Level; Set by the processor when a Reset, Soft Reset, NMI or Cache Error exception are taken.			
		Encoding Meaning			
		0 Normal level			
		1 Error level			
ERL	2	<ul> <li>When ERL is set:</li> <li>The processor is running in kernel mode</li> <li>Hardware and software interrupts are disabled</li> <li>The ERET instruction will use the return address held in ErrorEPC instead of EPC</li> <li>The lower 2<sup>29</sup> bytes of kuseg are treated as an unmapped and uncached region. See Section 4.7, "Address Translation for the kuseg Segment when StatusERL = 1" on page 16. This allows main memory to be accessed in the presence of cache errors. The operation of the processor is UNDEFINED if the ERL bit is set while the processor is executing instructions from kuseg.</li> </ul>	R/W	1	Required

Field	ds				Read/			
Name	Bits		Description		Write	Reset State	Compliance	
		Exception Le other than Re exception are	evel; Set by the processor when any exce eset, Soft Reset, NMI or Cache Error e taken.	ption				
		Encoding	Meaning	7				
		0	Normal level					
		1	Exception level					
EXL	1	When EXL • The proce • Hardware • TLB Refil instead of • EPC, Cau Release 2 if another	is set: ssor is running in Kernel Mode and software interrupts are disabled. l exceptions use the general exception v the TLB Refill vector. se <sub>BD</sub> and SRSCtl (implementations of of the Architecture only) will not be upp exception is taken	ector dated	R/W	Undefined	Required	
		Interrupt Ena and hardwar	able: Acts as the master enable for softw e interrupts: Meaning	vare				
IE	0	0	Interrupts are disabled	1	R/W	Undefined	Required	
	-	1	Interrupts are enabled	1				
		In Release 2 separately vi	of the Architecture, this bit may be mod a the DI and EI instructions.	lified				

I

# 8.18 IntCtl Register (CP0 Register 12, Select 1)

Compliance Level: Required (Release 2).

The IntCtl register controls the expanded interrupt capability added in Release 2 of the Architecture, including vectored interrupts and support for an external interrupt controller. This register does not exist in implementations of Release 1 of the Architecture.

Figure 8-16 shows the format of the IntCtl register; Table 8-20 describes the IntCtl register fields.

### Figure 8-16 IntCtl Register Format

31 29	28 26	25 10	9 5	4 0
IPTI	IPPCI	0 00 0000 0000 0000 00	VS	0

Fiel	Fields						Read/	Reset			
Name	Bits			Descri	ption		Write	State	Compliance		
		For Inter modes, ti Timer In software for a pot	rrupt Com his field sp nterrupt re to determ cential inte	patibility pecifies th quest is n tine whet errupt.	and Vectored Internet	errupt nich the s ause <sub>TI</sub>	e				
		] ]	Encoding	IP bit	Hardware Interrunt Source						
		-	2	2	HW0						
			3 3 HW1		Preset or						
IPTI	3129		4	4	HW2	-	R	Externally Set	Required		
			5	5	HW3						
			6	6	HW4						
			7	7	HW5						
			te of this f Interrupt ented and er is expect rrupt mod	field is UI Controlle enabled. ' eted to pro e.	NPREDICTABL er Mode is both The external inter ovide this informa	E if rupt tion for					

### Table 8-20 IntCtl Register Field Descriptions

Fiel	Fields						Read/	Pasat		
Name	Bits			Descri	ption		Write	State	Compliance	
		For Intern modes, thi Performar and allows Cause <sub>PCI</sub>	upt Con is field s nce Cou s softwa for a po	npatibility pecifies th nter Intern re to deter tential int	and Vectored Inte ne IP number to wh rupt request is mer mine whether to co errupt.	errupt iich the ged, onsider				
		E	ncoding	IP bit	Hardware Interrupt Source					
			2	2	HW0					
			3	3	HW1					
			4	4	HW2			Preset or	Optional (Performance	
IPPCI	PCI 2826		5	5	HW3		R	Externally Set	Counters	
			6	6	HW4			ber	Implemented)	
			7	7	HW5					
		External I implemen controller that interr If perform (Config1 <sub>P</sub>	The value of this field is UN External Interrupt Controlle implemented and enabled. T controller is expected to pro that interrupt mode. If performance counters are (Config1 <sub>PC</sub> = 0), this field re		er Mode is both The external interr ovide this informat e not implemented returns zero on rea	upt ion for d.				
0	2510	Must be w	vritten a	s zero; ret	urns zero on read.		0	0	Reserved	
		Vector Spi implemen Config3 <sub>VI</sub> between v	acing. If ted (as c EIC), this ectored	f vectored denoted by s field spe interrupts ng Betwee	interrupts are y Config3 <sub>VInt</sub> or cifies the spacing s. <b>n Spacing Betw</b>	een				
			Vec	tors (hex)	Vectors (decin	nal)				
		16#00	1	16#000	0					
		16#01	1	16#020	32					
VS	9.5	16#02	1	16#040	64		R/W	0	Optional	
		16#04	1	16#080	128		10/11		Optional	
		16#08	1	16#100	256					
		16#10	1	16#200	512					
		All other y processor written to If neither implemen 0), this fie	values a is <b>UND</b> this fiel EIC inte ted (Con ld is ign	re reserve DEFINED d. errupt moo nfig3 <sub>VEIC</sub> nored on v	d. The operation of if a reserved valu de nor VI mode ar = 0 and Config3 <sub>V</sub> vrite and reads as a	of the e is $e_{\rm TINT} = 2$ zero.				
0	40	Must be w	Aust be written as zero; returns zero on read.				0	0	Reserved	

# Table 8-20 IntCtl Register Field Descriptions

I

# 8.19 SRSCtl Register (CP0 Register 12, Select 2)

Compliance Level: Required (Release 2).

The *SRSCtl* register controls the operation of GPR shadow sets in the processor. This register does not exist in implementations of the architecture prior to Release 2.

Figure 8-17 shows the format of the SRSCtl register; Table 8-21 describes the SRSCtl register fields.

	Figure 8-17 SRSCtl Register Format																	
31 30	29		26	25	22	21 18	17	16 15	12	11 10	9		6	5	4	3		0
0 00		HSS		0 00 00		EICSS	0 00	)	ESS	0 00		PSS		( 0	) 0		CSS	

### **Table 8-21 SRSCtl Register Field Descriptions**

Fields			Read/	Reset	
Name	Bits	Description	Write	State	Compliance
0	3130	Must be written as zeros; returns zero on read.	0	0	Reserved
HSS	2926	Highest Shadow Set. This field contains the highest shadow set number that is implemented by this processor. A value of zero in this field indicates that only the normal GPRs are implemented. The value in this field also represents the highest value that can be written to the ESS, EICSS, PSS, and CSS fields of this register, or to any of the fields of the <i>SRSMap</i> register. The operation of the processor is <b>UNDEFINED</b> if a value larger than the one in this field is written to any of these other values.	R	Preset	Required
0	2522	Must be written as zeros; returns zero on read.	0	0	Reserved
EICSS	2118	<ul> <li>EIC interrupt mode shadow set. If Config3<sub>VEIC</sub> is 1 (EIC interrupt mode is enabled), this field is loaded from the external interrupt controller for each interrupt request and is used in place of the <i>SRSMap</i> register to select the current shadow set for the interrupt.</li> <li>See Section 5.1.1.3, "External Interrupt Controller Mode" on page 29 for a discussion of EIC interrupt mode. If Config3<sub>VEIC</sub> is 0, this field must be written as zero, and returns zero on read.</li> </ul>	R	Undefined	Required (EIC interrupt mode only)
0	1716	Must be written as zeros; returns zero on read.	0	0	Reserved
ESS	1512	Exception Shadow Set. This field specifies the shadow set to use on entry to Kernel Mode caused by any exception other than a vectored interrupt. The operation of the processor is <b>UNDEFINED</b> if software writes a value into this field that is greater than the value in the HSS field.	R/W	0	Required
0	1110	Must be written as zeros; returns zero on read.	0	0	Reserved

Field	ds		Read/	Rosot	
Name	Bits	Description	Write	State	Compliance
Dec	0.6	Previous Shadow Set. If GPR shadow registers are implemented, and with the exclusions noted in the next paragraph, this field is copied from the CSS field when an exception or interrupt occurs. An ERET instruction copies this value back into the CSS field if Status <sub>BEV</sub> = 0.	DAV	0	D
PSS	96	This field is not updated on any exception which sets Status <sub>ERL</sub> to 1 (i.e., NMI or cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with Status <sub>EXL</sub> = 1, or Status <sub>BEV</sub> = 1.	R/W	0	Required
		The operation of the processor is <b>UNDEFINED</b> if software writes a value into this field that is greater than the value in the HSS field.			
0	54	Must be written as zeros; returns zero on read.	0	0	Reserved
CSS	30	Current Shadow Set. If GPR shadow registers are implemented, this field is the number of the current GPR set. With the exclusions noted in the next paragraph, this field is updated with a new value on any interrupt or exception, and restored from the PSS field on an ERET. Table 8-22 describes the various sources from which the CSS field is updated on an exception or interrupt. This field is not updated on any exception which sets Status <sub>ERL</sub> to 1 (i.e., NMI or cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with Status <sub>EXL</sub> = 1, or Status <sub>BEV</sub> = 1. Neither is it updated on an ERET with Status <sub>ERL</sub> = 1 or Status <sub>BEV</sub> = 1. The value of CSS can be changed directly by software only by writing the PSS field and executing an ERET instruction.	R	0	Required

### Table 8-21 SRSCtl Register Field Descriptions

### Table 8-22 Sources for new $\mbox{SRSCtl}_{\mbox{CSS}}$ on an Exception or Interrupt

Exception Type	Condition	SRSCtl <sub>CSS</sub> Source	Comment
Exception	All	SRSCtl <sub>ESS</sub>	
Non-Vectored Interrupt	$Cause_{IV} = 0$	SRSCtl <sub>ESS</sub>	Treat as exception
Vectored Interrupt	$Cause_{IV} = 1 \text{ and} \\ Config3_{VEIC} = 0 \text{ and} \\ Config3_{VInt} = 1$	SRSMap <sub>VectNum</sub> ×4+3VectNum×4	Source is internal map register
Vectored EIC Interrupt	Cause <sub>IV</sub> = 1 and Config $3_{VEIC}$ = 1	SRSCtl <sub>EICSS</sub>	Source is external interrupt controller.

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

## 8.20 SRSMap Register (CP0 Register 12, Select 3)

**Compliance Level:** *Required* in Release 2 of the Architecture if Additional Shadow Sets and Vectored Interrupt Mode are Implemented

The *SRSMap* register contains 8 4-bit fields that provide the mapping from an vector number to the shadow set number to use when servicing such an interrupt. The values from this register are not used for a non-interrupt exception, or a non-vectored interrupt (Cause<sub>IV</sub> = 0 or IntCtl<sub>VS</sub> = 0). In such cases, the shadow set number comes from SRSCtl<sub>ESS</sub>.

If SRSCtl<sub>HSS</sub> is zero, the results of a software read or write of this register are UNPREDICTABLE.

The operation of the processor is **UNDEFINED** if a value is written to any field in this register that is greater than the value of SRSCtl<sub>HSS</sub>.

The *SRSMap* register contains the shadow register set numbers for vector numbers 7..0. The same shadow set number can be established for multiple interrupt vectors, creating a many-to-one mapping from a vector to a single shadow register set number.

Figure 8-18 shows the format of the SRSMap register; Table 8-23 describes the SRSMap register fields.

#### Figure 8-18 SRSMap Register Format

31	28 27	24	23	20	19	1	16	15		12	11		8	7		4	3		0
SSV7	SSV6		SSV	5		SSV4			SSV3			SSV2			SSV1			SSV0	

#### Table 8-23 SRSMap Register Field Descriptions

Fiel	ds		Read/		
Name	Bits	Description	Write	Reset State	Compliance
SSV7	3128	Shadow register set number for Vector Number 7	R/W	0	Required
SSV6	2724	Shadow register set number for Vector Number 6	R/W	0	Required
SSV5	2320	Shadow register set number for Vector Number 5	R/W	0	Required
SSV4	1916	Shadow register set number for Vector Number 4	R/W	0	Required
SSV3	1512	Shadow register set number for Vector Number 3	R/W	0	Required
SSV2	118	Shadow register set number for Vector Number 2	R/W	0	Required
SSV1	74	Shadow register set number for Vector Number 1	R/W	0	Required
SSV0	30	Shadow register set number for Vector Number 0	R/W	0	Required

# 8.21 Cause Register (CP0 Register 13, Select 0)

### **Compliance Level:** *Required.*

The *Cause* register primarily describes the cause of the most recent exception. In addition, fields also control software interrupt requests and the vector through which interrupts are dispatched. With the exception of the  $IP_{1..0}$ , DC, IV, and WP fields, all fields in the Cause register are read-only. Release 2 of the Architecture added optional support for an External Interrupt Controller (EIC) interrupt mode, in which  $IP_{7..2}$  are interpreted as the Requested Interrupt Priority Level (RIPL).

Figure 8-19 shows the format of the Cause register; Table 8-24 describes the Cause register fields.

Figure 8-19 Cause Register Format																				
31 30	29 28	8 27	26	25	24	23	22	21	16	15		10	9	8	7	6		2	1	0
BD TI	CE	DC	PCI	(	)	IV	WP	0			IP7IP2		IP1.	.IP0	0		Exc Code		C	)
											RIPL									

Fields				Read/			
Name	Bits		Description	Write	Reset State	Compliance	
		Indicates v a branch d	whether the last exception taken occurred in elay slot:				
		Encoding	Meaning				
BD	31	0	Not in delay slot	R	Undefined	Required	
DD	51	1	In delay slot	K	Chaemiea	Required	
		The proces when the	ssor updates BD only if Status <sub>EXL</sub> was zero exception occurred.				
		the Archit interrupt is interrupt t	rrupt. In an implementation of Release 2 of ecture, this bit denotes whether a timer s pending (analogous to the IP bits for other ypes):				
		Encoding	Meaning			Required	
TI	30	0	No timer interrupt is pending	R	Undefined	(Release 2)	
		1	Timer interrupt is pending				
		In an impl Architectu returns zer	ementation of Release 1 of the re, this bit must be written as zero and o on read.				
CE	2928	Coprocess Coprocess is loaded I UNPRED Coprocess	or unit number referenced when a or Unusable exception is taken. This field by hardware on every exception, but is <b>ICTABLE</b> for all exceptions except for or Unusable.	R	Undefined	Required	

### **Table 8-24 Cause Register Field Descriptions**

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

# Table 8-24 Cause Register Field Descriptions

Fiel	ds		Dood/		
Name	Bits	Description	Write	Reset State	Compliance
		Disable <i>Count</i> register. In some power-sensitive applications, the <i>Count</i> register is not used but may still be the source of some noticeable power dissipation. This bit allows the <i>Count</i> register to be stopped in such situations.			
DC	27	Encoding Meaning	R/W	0	Required
		Disable counting of <i>Count</i> register     Disable counting of <i>Count</i> register			(Release 2)
		In an implementation of Release 1 of the Architecture, this bit must be written as zero, and returns zero on read.			
		Performance Counter Interrupt. In an implementation of Release 2 of the Architecture, this bit denotes whether a performance counter interrupt is pending (analogous to the IP bits for other interrupt types):			
		Encoding Meaning			Required (Release 2 and
PCI	26	0 No timer interrupt is pending	R	Undefined	performance
		1 Timer interrupt is pending			implemented)
		In an implementation of Release 1 of the Architecture, or if performance counters are not implemented (Config $1_{PC} = 0$ ), this bit must be written as zero and returns zero on read.			
		Indicates whether an interrupt exception uses the general exception vector or a special interrupt vector:			
		Encoding Meaning			
		0 Use the general exception vector (16#180)			
IV	23	1 Use the special interrupt vector (16#200)	R/W	Undefined	Required
		In implementations of Release 2 of the architecture, if the Cause <sub>IV</sub> is 1 and Status <sub>BEV</sub> is 0, the special interrupt vector represents the base of the vectored interrupt table.			
		Indicates that a watch exception was deferred because $Status_{EXL}$ or $Status_{ERL}$ were a one at the time the watch exception was detected. This bit both indicates that the watch exception was deferred, and causes the exception to be initiated once $Status_{EXL}$ and $Status_{ERL}$ are both zero. As such, software must clear this bit as part of the watch exception handler to prevent a watch exception loop.			Promined if
WP	22	Software should not write a 1 to this bit when its value is a 0, thereby causing a 0-to-1 transition. If such a transition is caused by software, it is <b>UNPREDICTABLE</b> whether hardware ignores the write, accepts the write with no side effects, or accepts the write and initiates a watch exception once Status <sub>EXL</sub> and Status <sub>ERL</sub> are both zero.	R/W	Undefined	watch registers are implemented
		If watch registers are not implemented, this bit must be ignored on write and read as zero.			

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

Fields						Read/		
Name	Bits			Description		Write	Reset State	Compliance
		Indicates	an inter	rrupt is pending:				
		Bit	Name	Meaning				
		15	IP7	Hardware interrupt 5				
		14	IP6	Hardware interrupt 4				
		13	IP5	Hardware interrupt 3				
		12	IP4	Hardware interrupt 2				
		11	IP3	Hardware interrupt 1				
		10	IP2	Hardware interrupt 0				
	1510	In implem timer and combined hardware In implem in which H (Config3 <sub>V</sub> interrupts implemen interrupt. (Config3 <sub>w</sub> meaning a described	entation perform in an in interrup entation EIC inter FEIC = ( are con- tation-out If EIC If EIC mod are below.	ons of Release 1 of the Architec mance counter interrupts are mplementation-dependent way pt 5. ons of Release 2 of the Architec errupt mode is not enabled )), timer and performance coun nbined in an dependent way with any hardw interrupt mode is enabled 1), these bits take on a differen interpreted as the RIPL field,	eture, with cture nter vare t			
RIPL	1510	Requested In implem in which E = 1), this i requested interrupt i If EIC inte 0), these b interpreted	l Interru eentatic EIC inte field is interru s reque errupt r bits take d as the	upt Priority Level. ons of Release 2 of the Archite. rrupt mode is enabled (Config: the encoded (063) value of th pt. A value of zero indicates th sted. node is not enabled (Config3 <sub>V</sub> e on a different meaning and a e IP7IP2 bits, described above	R	Undefined	Optional (Release 2 and EIC interrupt mode only)	
		Controls t	he requ	lest for software interrupts:				
		Bit	Name	Meaning				
		9	IP1	Request software interrupt 1				
IP1IP0	98	8	IP0	Request software interrupt 0		R/W	Undefined	Required
		An impler which also these bits prioritizat	mentati 5 imple to the e ion wit	on of Release 2 of the Archite ments EIC interrupt mode exp external interrupt controller for h other interrupt sources.	itecture xports for			
ExcCode	62	Exception	code -	see Table 8-25		R	Undefined	Required
0	2524, 2116, 7, 10	Must be w	ritten a	as zero; returns zero on read.		0	0	Reserved

# Table 8-24 Cause Register Field Descriptions

Exception Code Value			
Decimal	Hexadecimal	Mnemonic	Description
0	16#00	Int	Interrupt
1	16#01	Mod	TLB modification exception
2	16#02	TLBL	TLB exception (load or instruction fetch)
3	16#03	TLBS	TLB exception (store)
4	16#04	AdEL	Address error exception (load or instruction fetch)
5	16#05	AdES	Address error exception (store)
6	16#06	IBE	Bus error exception (instruction fetch)
7	16#07	DBE	Bus error exception (data reference: load or store)
8	16#08	Sys	Syscall exception
9	16#09	Вр	Breakpoint exception. If EJTAG is implemented and an SDBBP instruction is executed while the processor is running in EJTAG Debug Mode, this value is written to the Debug <sub>DExcCode</sub> field to denote an SDBBP in Debug Mode.
10	16#0a	RI	Reserved instruction exception
11	16#0b	CpU	Coprocessor Unusable exception
12	16#0c	Ov	Arithmetic Overflow exception
13	16#0d	Tr	Trap exception
14	16#0e	-	Reserved
15	16#0f	FPE	Floating point exception
16-17	16#10-16#11	-	Available for implementation dependent use
18	16#12	C2E	Reserved for precise Coprocessor 2 exceptions
19-21	16#13-16#15	-	Reserved
22	16#16	MDMX	MDMX Unusable Exception.
23	16#17	WATCH	Reference to WatchHi/WatchLo address
24	16#18	MCheck	Machine check
25-29	16#19-16#1d	-	Reserved
30	16#1e	CacheErr	Cache error. In normal mode, a cache error exception has a dedicated vector and the Cause register is not updated. If EJTAG is implemented and a cache error occurs while in Debug Mode, this code is writen to the Debug <sub>DExcCode</sub> field to indicate that re-entry to Debug Mode was caused by a cache error.
31	16#1f	-	Reserved

# Table 8-25 Cause Register ExcCode Field

### 8.22 Exception Program Counter (CP0 Register 14, Select 0)

#### Compliance Level: Required.

The *Exception Program Counter (EPC)* is a read/write register that contains the address at which processing resumes after an exception has been serviced. All bits of the *EPC* register are significant and must be writable.

For synchronous (precise) exceptions, EPC contains either:

- the virtual address of the instruction that was the direct cause of the exception, or
- the virtual address of the immediately preceding branch or jump instruction, when the exception causing instruction is in a branch delay slot, and the *Branch Delay* bit in the *Cause* register is set.

For asynchronous (imprecise) exceptions, EPC contains the address of the instruction at which to resume execution.

The processor does not write to the *EPC* register when the EXL bit in the *Status* register is set to one.

Figure 8-20 shows the format of the EPC register; Table 8-26 describes the EPC register fields.

#### Figure 8-20 EPC Register Format

31	0
	EPC

#### **Table 8-26 EPC Register Field Descriptions**

Fiel	ds		Read/			
Name	Bits	Description	Write	Reset State	Compliance	
EPC	310	Exception Program Counter	R/W	Undefined	Required	

### 8.22.1 Special Handling of the EPC Register in Processors That Implement the MIPS16e ASE

In processors that implement the MIPS16e ASE, a read of the *EPC* register (via MFC0) returns the following value in the destination GPR:

 $GPR[rt] \leftarrow RestartPC_{31..1} || ISAMode$ 

That is, the upper 31 bits of the restart PC are combined with the ISA Mode bit and written to the GPR.

Similarly, a write to the *EPC* register (via MTC0) takes the value from the GPR and distributes that value to the restart PC and the *ISA Mode* bit, as follows

 $\begin{aligned} \text{RestartPC} &\leftarrow \text{GPR[rt]}_{31..1} || & 0\\ \text{ISAMode} &\leftarrow \text{GPR[rt]}_{0} \end{aligned}$ 

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the restart PC, and the lower bit of the restart PC is cleared. The *ISA Mode* bit is loaded from the lower bit of the GPR.

## 8.23 Processor Identification (CP0 Register 15, Select 0)

### Compliance Level: Required.

The *Processor Identification (PRId)* register is a 32 bit read-only register that contains information identifying the manufacturer, manufacturer options, processor identification and revision level of the processor. Figure 8-21 shows the format of the *PRId* register; Table 8-27 describes the *PRId* register fields.

### Figure 8-21 PRId Register Format

31	24	23 16	15 8	7 0
	Company Options	Company ID	Processor ID	Revision

Fields				Read/		
Name	Bits		Description	Write	Reset State	Compliance
Company Options	3124	Available processor value in th If this field	to the designer or manufacturer of the for company-dependent options. The is field is not specified by the architecture. d is not implemented, it must read as zero.	R	Preset	Optional
Company ID	2316	Identifies manufactures of tware of processor ISA by chemical the process Architecture Company when a Mencodings Encoding 0 1 2-255	the company that designed or red the processor. can distinguish a MIPS32 or MIPS64 from one implementing an earlier MIPS ecking this field for zero. If it is non-zero sor implements the MIPS32 or MIPS64 re. IDs are assigned by MIPS Technologies IPS32 or MIPS64 license is acquired. The in this field are: $\label{eq:meansature} \begin{tabular}{lllllllllllllllllllllllllllllllllll$	R	Preset	Required
Processor ID	158	Identifies software to implemen qualified b The comb Processorl to each pro	the type of processor. This field allows to distinguish between various processor tations within a single company, and is by the CompanyID field, described above. ination of the CompanyID and ID fields creates a unique number assigned processor implementation.	R	Preset	Required
Revision	70	Specifies t field allow revision an this field in	he revision number of the processor. This is software to distinguish between one ad another of the same processor type. If is not implemented, it must read as zero.	R	Preset	Optional

### Table 8-27 PRId Register Field Descriptions

Software should not use the fields of this register to infer configuration information about the processor. Rather, the configuration registers should be used to determine the capabilities of the processor. Programmers who identify cases in which the configuration registers are not sufficient, requiring them to revert to check on the *PRId* register value, should send email to architecture@mips.com, reporting the specific case.

# 8.24 EBase Register (CP0 Register 15, Select 1)

Compliance Level: Required (Release 2).

The *EBase* register is a read/write register containing the base address of the exception vectors used when  $Status_{BEV}$  equals 0, and a read-only CPU number value that may be used by software to distinguish different processors in a multi-processor system.

The *EBase* register provides the ability for software to identify the specific processor within a multi-processor system, and allows the exception vectors for each processor to be different, especially in systems composed of heterogeneous processors. Bits 31..12 of the *EBase* register are concatenated with zeros to form the base of the exception vectors when Status<sub>BEV</sub> is 0. The exception vector base address comes from the fixed defaults (see Section 5.2.1, "Exception Vector Locations" on page 32) when Status<sub>BEV</sub> is 1, or for any EJTAG Debug exception. The reset state of bits 31..12 of the *EBase* register initialize the exception base register to 16#8000.0000, providing backward compatibility with Release 1 implementations.

Bits 31..30 of the *EBase* Register are fixed with the value 2#10, and the addition of the base address and the exception offset is done inhibiting a carry between bit 29 and bit 30 of the final exception address. The combination of these two restrictions forces the final exception address to be in the kseg0 or kseg1 unmapped virtual address segments. For cache error exceptions, bit 29 is forced to a 1 in the ultimate exception base address so that this exception always runs in the kseg1 unmapped, uncached virtual address segment.

If the value of the exception base register is to be changed, this must be done with  $Status_{BEV}$  equal 1. The operation of the processor is **UNDEFINED** if the Exception Base field is written with a different value when  $Status_{BEV}$  is 0.

Figure 8-22 shows the format of the EBase Register; Table 8-28 describes the EBase register fields.

### Figure 8-22 EBase Register Format

31 30 29			2 11 10	9	0
1	0	Exception Base	0 0	CPUNum	

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
1	31	This bit is ignored on write and returns one on read.	R	1	Required
0	30	This bit is ignored on write and returns zero on read.	R	0	Required
Exception Base	2912	In conjunction with bits $3130$ , this field specifies the base address of the exception vectors when $Status_{BEV}$ is zero.	R/W	0	Required
0	1110	Must be written as zero; returns zero on read.	0	0	Reserved
CPUNum	90	This field specifies the number of the CPU in a multi-processor system and can be used by software to distinguish a particular processor from the others. The value in this field is set by inputs to the processor hardware when the processor is implemented in the system environment. In a single processor system, this value should be set to zero.	R	Preset or Externally Set	Required

### Table 8-28 EBase Register Field Descriptions

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

#### **Programming Note:**

Software must set  $EBase_{15..12}$  to zero in all bit positions less than or equal to the most significant bit in the vector offset. This situation can only occur when a vector offset greater than 16#FFF is generated when an interrupt occurs with VI or EIC interrupt mode enabled. The operation of the processor is **UNDEFINED** if this condition is not met. Table 8-29 shows the conditions under which each EBase bit must be set to zero. VN represents the interrupt vector number as described in Table 5-4 on page 32 and the bit must be set to zero if any of the relationships in the row are true. No EBase bits must be set to zero if the interrupt vector spacing is 32 (or zero) bytes.

	Interrupt Vector Spacing in Bytes (IntCtl <sub>VS</sub> <sup>1</sup> )									
EBase bit	se bit 32 64 128 256				512					
15		None	None	None	VN ≥ 63					
14	Nona	None	Νονε	VN ≥ 62	VN ≥ 31					
13	None	Νονε	VN ≥ 60	VN ≥ 30	VN ≥ 15					
12		VN ≥ 56	VN ≥ 28	VN ≥ 14	$VN \ge 7$					

Table 8-29 Conditions Under Which EBase15..12 Must Be Zero

1. See Table 8-20 on page 82

### 8.25 Configuration Register (CP0 Register 16, Select 0)

### Compliance Level: Required.

The *Config* register specifies various configuration and capabilities information. Most of the fields in the *Config* register are initialized by hardware during the Reset Exception process, or are constant. One field, K0, must be initialized by software in the reset exception handler.

Figure 8-23 shows the format of the Config register; Table 8-30 describes the Config register fields.

### Figure 8-23 Config Register Format

31	30 16	15	14 13	12 10	9 7	6 4	3	2 0
M	Impl	BE	AT	AR	MT	0	VI	K0

Fields				Read/		
Name	Bits		Description	Write	Reset State	Compliance
М	31	Denotes that select field	at the Config1 register is implemented at a value of 1.	R	1	Required
Impl	30:16	This field is processor s this field	reserved for implementations. Refer to the pecification for the format and definition of		Undefined	Optional
BE	15	Indicates th running: Encoding 0 1	Meaning Little endian Big endian	R	Preset or Externally Set	Required
AT	14:13	Architectur Encoding 0 1 2 3	e type implemented by the processor:         Meaning         MIPS32         MIPS64 with access only to 32-bit compatibility segments         MIPS64 with access to all address segments         Reserved	R	Preset	Required
AR	12:10	Architectur Encoding 0 1 2-7	e revision level: Meaning Release 1 Release 2 Reserved	R	Preset	Required

### **Table 8-30 Config Register Field Descriptions**

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

Fields				Read/		
Name	Bits		Description	Write	Reset State	Compliance
		MMU Type	::			
		Encoding	Meaning			
		0	None			
		1	Standard TLB		Preset	
MT	9:7	2	Standard BAT (see Section A.2 on page 131)	R		Required
		3	Standard fixed mapping (see Section A.1 on page 127)			
		4-7	Reserved			
0	6:4	Must be wr	itten as zero; returns zero on read.	0	0	Reserved
		Virtual inst and virtual	ruction cache (using both virtual indexing tags):			
VI	3	Encoding	Meaning	R	Preset	Required
	_	0	Instruction Cache is not virtual			1
		1	Instruction Cache is virtual			
КО	2:0	Kseg0 cohe for the ence	erency algorithm. See Table 8-8 on page 61 oding of this field.	R/W	Undefined	Optional

# Table 8-30 Config Register Field Descriptions

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

### 8.26 Configuration Register 1 (CP0 Register 16, Select 1)

### Compliance Level: Required.

The *Config1* register is an adjunct to the *Config* register and encodes additional capabilities information. All fields in the *Config1* register are read-only.

The Icache and Dcache configuration parameters include encodings for the number of sets per way, the line size, and the associativity. The total cache size for a cache is therefore:

Cache Size = Associativity \* Line Size \* Sets Per Way

If the line size is zero, there is no cache implemented.

Figure 8-24 shows the format of the Config1 register; Table 8-31 describes the Config1 register fields.

	Figure 8-24 Config1 Register Format																					
31	30	25	24	22	2 21	19	9 18	16	i 1:	5 1	13	12	10	9	7	6	5	4	3	2	1	0
М	MMU Size -	1		IS		IL		IA		DS		Ľ	DL		DA	C2	MD	PC	WR	CA	EP	FP

Fields				Read/			
Name	Bits		Description	Write	Reset State	Compliance	
М	31	This bit is r present. If t bit should r implemente	eserved to indicate that a <i>Config2</i> register is he <i>Config2</i> register is not implemented, this read as a 0. If the <i>Config2</i> register is ed, this bit should read as a 1.	R	Preset	Required	
MMU Size - 1	3025	Number of through 63 entries. The a value of '	entries in the TLB minus one. The values 0 is this field correspond to 1 to 64 TLB e value zero is implied by Config <sub>MT</sub> having none'.	R	Preset	Required	
		Icache sets	per way:			Required	
		Encoding	Meaning				
		0	64		Preset		
		1	128	R			
IS	24.22	2	256				
15	24.22	3	512			Required	
		4	1024				
		5	2048				
		6	4096				
		7	Reserved				

Fields					Road/		
Name	Bits		Description		Write	Reset State	Compliance
		Icache line	size:				
		Encoding	Encoding Meaning				
		0	No Icache present				
		1	4 bytes				
п	21.10	2	8 bytes		р	Dragat	Dequired
IL	21.19	3	16 bytes		К	Fleset	Kequileu
		4	32 bytes				
		5	64 bytes				
		6	128 bytes				
		7	Reserved				
		Icache asso	ciativity:				
		Encoding	Meaning	1			
	18:16	0	Direct mapped	- R			
		1	2-way				
TA		2	3-way		D		D . I
IA		3	4-way		ĸ	Preset	Required
		4	5-way				
		5	6-way				
		6	7-way				
		7	8-way				
		Dcache set	s per way:				
		Encoding	Meaning	]			l
		0	64				
		1	128			Preset	
DS	15:13	2	256	1	R		Required
25	15:15	3	512		R		riequireu
		4	1024				
		5	2048				
		6	4096				
		7	Reserved	]			

Fields				Dood/			
Name	Bits		Description	Write	Reset State	Compliance	
		Dcache line	e size:				
		Encoding	Meaning				
		0	No Dcache present				
		1	4 bytes				
Ы	12.10	2	8 bytes	р	Dreast	Dequired	
DL	12:10	3	16 bytes	ĸ	Preset	Required	
		4	32 bytes				
		5	64 bytes				
		6	128 bytes				
		7	Reserved				
		Dcache ass	ociativity:				
		Encoding	Meaning				
		0	Direct mapped		Preset		
	9:7	1	2-way	R			
		2	3-way				
DA		3	4-way			Required	
		4	5-way				
		5	6-way				
		6	7-way				
		7	8-way				
		Coprocesso	r 2 implemented:				
		Encoding	Meaning				
		0	No coprocessor 2 implemented				
C2	6	1	Coprocessor 2 implements				
		This bit ind support for is attached.	cates not only that the processor contains Coprocessor 2, but that such a coprocessor				
		Used to der MIPS64 pr	note MDMX ASE implemented on a occessor. Not used on a MIPS32 processor.		0		
MD	5	This bit ind support for is attached.	icates not only that the processor contains MDMX, but that such a processing element	R		Required	
		Performanc	e Counter registers implemented:				
		Encoding	Meaning				
PC	4	0	No performance counter registers implemented	R	Preset	Required	
		1	Performance counter registers implemented				

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
		Watch registers implemented:			
WD	2	Encoding Meaning	D	Durant	D a suciona d
WK	5	0 No watch registers implemented	К	Preset	Required
		1 Watch registers implemented			
		Code compression (MIPS16e) implemented:			
		Encoding Meaning		Preset	Required
CA	2	0 MIPS16e not implemented	R		
		1 MIPS16e implemented			
EP	1	EJTAG implemented:         Encoding       Meaning         0       No EJTAG implemented         1       EJTAG implemented	R	Preset	Required
		FPU implemented:			
		Encoding Meaning		Preset	
		0 No FPU implemented			
		1 FPU implemented			
FP	0	This bit indicates not only that the processor contains support for a floating point unit, but that such a unit is attached. If an FPU is implemented, the capabilities of the FPU can be read from the capability bits in the <i>FIR</i> CP1 register.	R		Required
### 8.27 Configuration Register 2 (CP0 Register 16, Select 2)

**Compliance Level:** Required if a level 2 or level 3 cache is implemented, or if the Config3 register is required; Optional otherwise.

The Config2 register encodes level 2 and level 3 cache configurations.

Figure 8-25 shows the format of the Config2 register; Table 8-32 describes the Config2 register fields.

#### Figure 8-25 Config2 Register Format

31	30 2	8 27	2	4 23	2	0 19		16	15		12	11		8	7		4	3		0
М	TU		TS		TL		TA			SU			SS			SL			SA	

Fiel	lds				Read/		
Fields       Name     Bits       M     31       TU     30:28		De	Write	Reset State	Compliance		
М	31	This bit is reserved to inc present. If the Config3 re bit should read as a 0. If implemented, this bit sh	dicate that a Co egister is not in the Config3 ro ould read as a	onfig3 register is nplemented, this egister is 1.	R	Preset	Required
TU	30:28	Implementation-specific bits. If this field is not in zero and be ignored on	tertiary cache nplemented it write.	control or status should read as	R/W	Preset	Optional
TS	27:24	Encoding           0           1           2           3           4           5           6           7           8-15	Sets Per Way           64           128           256           512           1024           2048           4096           8192           Reserved		R	Preset	Required

### Table 8-32 Config2 Register Field Descriptions

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

I

Fields			Dood/		
Name	Bits	Description	Write	Reset State	Compliance
TL	23:20	Tertiary cache line size:EncodingLine Size0No cache present1428316432564612872568-15Reserved	R	Preset	Required
TA	19:16	Tertiary cache associativity:EncodingAssociativity0Direct Mapped122334455667788-15Reserved	R	Preset	Required
SU	15:12	Implementation-specific secondary cache c status bits. If this field is not implemented i read as zero and be ignored on write.	control or it should R/W	Preset	Optional
SS	11:8	Encoding         Sets Per Way           0         64           1         128           2         256           3         512           4         1024           5         2048           6         4096           7         8192           8-15         Reserved	R	Preset	Required

### Table 8-32 Config2 Register Field Descriptions

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

I

I

I

Fiel	ds			Read/			
Name	Bits	Des	scription	Write	Reset State	Compliance	
		Secondary cache line siz					
		Encoding	Line Size				
		0	No cache present				
		1	1 4				
SI	7.4	2	8	D	Drasat	Paquirad	
SL	/.4	3	16	K	Fleset	Required	
		4	32				
		5	64				
		6	128				
		7	256				
		8-15	Reserved				
		Secondary cache associa	utivity:				
		Encoding	Associativity				
		0	Direct Mapped				
		1	2				
SA	3.0	2	3	R	Preset	Required	
571	5.0	3	4	K	Treset	Required	
		4	5				
		5	6				
		6	7				
		7	8				
		8-15	Reserved				

### Table 8-32 Config2 Register Field Descriptions

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

I

I

I

# 8.28 Configuration Register 3 (CP0 Register 16, Select 3)

**Compliance Level:** *Required if any optional feature described by this register is implemented: Release 2 of the Architecture, the SmartMIPS ASE, or trace logic; optional otherwise.* 

The Config3 register encodes additional capabilities. All fields in the Config3 register are read-only.

Figure 8-26 shows the format of the Config3 register; Table 8-33 describes the Config3 register fields.

#### Figure 8-26 Config3 Register Format

31	30	7	6	5	4	3	2	1	0
M	0 000 0000 0000 0000 0000 0000	LPA	VEIC	VInt	SP	0		SM	TL

Fields			Read/		
Name Bits		Description	Write	Reset State	Compliance
М	31	This bit is reserved to indicate that a Config4 register is present. With the current architectural definition, this bit should always read as a 0.	R	Preset	Required
0	30:8,3:2	Must be written as zeros; returns zeros on read	0	0	Reserved
LPA	7	Denotes the presence of support for large physical addresses on MIPS64 processors. Not used by MIPS32 processors and returns zero on read. For implementations of Release 1 of the Architecture, this bit returns zero on read.	R	Preset	Required (Release 2 Only)
VEIC	6	Support for an external interrupt controller is implemented.         Encoding       Meaning         0       Support for EIC interrupt mode is not implemented         1       Support for EIC interrupt mode is implemented         For implementations of Release 1 of the Architecture, this bit returns zero on read.         This bit indicates not only that the processor contains support for an external interrupt controller, but that such a controller is attached.	R	Preset	Required (Release 2 Only)
VInt	5	Vectored interrupts implemented. This bit indicates whether vectored interrupts are implemented.         Encoding       Meaning         0       Vector interrupts are not implemented         1       Vectored interrupts are implemented         For implementations of Release 1 of the Architecture, this bit returns zero on read.	R	Preset	Required (Release 2 Only)

#### Table 8-33 Config3 Register Field Descriptions

Fields       Name     Bits				Read/				
			Description	Write	Reset State	Compliance		
		Small (1KBy PageGrain re	te) page support is implemented, ar gister exists					
		Encoding	Meaning	7			Required	
SP	4	0	Small page support is not implemented	1	R	Preset	(Release 2	
		1	Small page support is implemented				Only)	
		For implement this bit return	ntations of Release 1 of the Architer as zero on read.	cture,				
		SmartMIPS <sup>TI</sup> whether the S	MASE implemented. This bit indicates SmartMIPS ASE is implemented.	ates				
SM	1	Encoding	Meaning		R	Preset	Required	
		0	SmartMIPS ASE is not implemented					
		1	SmartMIPS ASE is implemented					
		Trace Logic i PC or data tra	implemented. This bit indicates whe ace is implemented.	ether				
TL	0	Encoding	Meaning		R	Preset	Required	
		0	Trace logic is not implemented					
		1	Trace logic is implemented					
				-				

### Table 8-33 Config3 Register Field Descriptions

### 8.29 Reserved for Implementations (CP0 Register 16, Selects 6 and 7)

Compliance Level: Optional: Implementation Dependent.

CP0 register 16, Selects 6 and 7 are reserved for implementation dependent use and is not defined by the architecture. In order to use CP0 register 16, Selects 6 and 7, it is not necessary to implement CP0 register 16, Selects 2 through 5 only to set the M bit in each of these registers. That is, if the *Config2* and *Config3* registers are not needed for the implementation, they need not be implemented just to provide the M bits.

### 8.30 Load Linked Address (CP0 Register 17, Select 0)

#### **Compliance Level:** *Optional.*

The *LLAddr* register contains relevant bits of the physical address read by the most recent Load Linked instruction. This register is implementation dependent and for diagnostic purposes only and serves no function during normal operation.

Figure 8-27 shows the format of the LLAddr register; Table 8-34 describes the LLAddr register fields.

#### Figure 8-27 LLAddr Register Format

31	0
	PAddr

### Table 8-34 LLAddr Register Field Descriptions

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
PAddr	310	This field encodes the physical address read by the most recent Load Linked instruction. The format of this register is implementation dependent, and an implementation may implement as many of the bits or format the address in any way that it finds convenient.	R	Undefined	Optional

### 8.31 WatchLo Register (CP0 Register 18)

#### Compliance Level: Optional.

The *WatchLo* and *WatchHi* registers together provide the interface to a watchpoint debug facility which initiates a watch exception if an instruction or data access matches the address specified in the registers. As such, they duplicate some functions of the EJTAG debug solution. Watch exceptions are taken only if the EXL and ERL bits are zero in the *Status* register. If either bit is a one, the WP bit is set in the *Cause* register, and the watch exception is deferred until both the EXL and ERL bits are zero.

An implementation may provide zero or more pairs of WatchLo and WatchHi registers, referencing them via the select field of the MTC0/MFC0 instructions, and each pair of Watch registers may be dedicated to a particular type of reference (e.g., instruction or data). Software may determine if at least one pair of *WatchLo* and *WatchHi* registers are implemented via the WR bit of the *Config1* register. See the discussion of the M bit in the *WatchHi* register description below.

The *WatchLo* register specifies the base virtual address and the type of reference (instruction fetch, load, store) to match. If a particular Watch register only supports a subset of the reference types, the unimplemented enables must be ignored on write and return zero on read. Software may determine which enables are supported by a particular Watch register pair by setting all three enables bits and reading them back to see which ones were actually set.

It is implementation dependent whether a data watch is triggered by a prefetch, CACHE, or SYNCI (Release 2 only) instruction whose address matches the Watch register address match conditions.

Figure 8-28 shows the format of the WatchLo register; Table 8-35 describes the WatchLo register fields.

#### Figure 8-28 WatchLo Register Format

31	3	2	1	0
VAddr		Ι	R	W

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
VAddr	313	This field specifies the virtual address to match. Note that this is a doubleword address, since bits [2:0] are used to control the type of match.	R/W	Undefined	Required
Ι	2	If this bit is one, watch exceptions are enabled for instruction fetches that match the address and are actually issued by the processor (speculative instructions never cause Watch exceptions). If this bit is not implemented, writes to it must be ignored, and reads must return zero.	R/W	0	Optional
R	1	If this bit is one, watch exceptions are enabled for loads that match the address. For the purposes of the MIPS16e PC-relative load instructions, the PC-relative reference is considered to be a data, rather than an instruction reference. That is, the watchpoint is triggered only if this bit is a 1. If this bit is not implemented, writes to it must be ignored, and reads must return zero.	R/W	0	Optional

#### Table 8-35 WatchLo Register Field Descriptions

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
W	0	If this bit is one, watch exceptions are enabled for stores that match the address. If this bit is not implemented, writes to it must be ignored, and reads must return zero.	R/W	0	Optional

# Table 8-35 WatchLo Register Field Descriptions

### 8.32 WatchHi Register (CP0 Register 19)

#### Compliance Level: Optional.

The *WatchLo* and *WatchHi* registers together provide the interface to a watchpoint debug facility which initiates a watch exception if an instruction or data access matches the address specified in the registers. As such, they duplicate some functions of the EJTAG debug solution. Watch exceptions are taken only if the EXL and ERL bits are zero in the *Status* register. If either bit is a one, the WP bit is set in the *Cause* register, and the watch exception is deferred until both the EXL and ERL bits are zero.

An implementation may provide zero or more pairs of WatchLo and WatchHi registers, referencing them via the select field of the MTC0/MFC0 instructions, and each pair of Watch registers may be dedicated to a particular type of reference (e.g., instruction or data). Software may determine if at least one pair of *WatchLo* and *WatchHi* registers are implemented via the WR bit of the *Config1* register. If the M bit is one in the *WatchHi* register reference with a select field of 'n', another WatchHi/WatchLo pair is implemented with a select field of 'n+1'.

The *WatchHi* register contains information that qualifies the virtual address specified in the *WatchLo* register: an ASID, a G(lobal) bit, an optional address mask, and three bits (I, R, and W) which denote the condition that caused the watch register to match. If the G bit is one, any virtual address reference that matches the specified address will cause a watch exception. If the G bit is a zero, only those virtual address references for which the ASID value in the *WatchHi* register matches the ASID value in the *EntryHi* register cause a watch exception. The optional mask field provides address masking to qualify the address specified in *WatchLo*.

The I, R, and W bits are set by the processor when the corresponding watch register condition is satisfied and indicate which watch register pair (if more than one is implemented) and which condition matched. When set by the processor, each of these bits remain set until cleared by software. All three bits are "write one to clear", such that software must write a one to the bit in order to clear its value. The typical way to do this is to write the value read from the *WatchHi* register back to *WatchHi*. In doing so, only those bits which were set when the register was read are cleared when the register is written back.

Figure 8-29 shows the format of the WatchHi register; Table 8-36 describes the WatchHi register fields.

				Figure 8-29 Wa	itchHi Regist	ter Format				
31	30	29	24	23 16	15 12	11	3	2	1	0
Μ	G	0		ASID	0	Mask		Ι	R	W

#### Table 8-36 WatchHi Register Field Descriptions

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
М	31	If this bit is one, another pair of <i>WatchHi/WatchLo</i> registers is implemented at a MTC0 or MFC0 select field value of $n+1$ ?	R	Preset	Required
G	30	If this bit is one, any address that matches that specified in the <i>WatchLo</i> register will cause a watch exception. If this bit is zero, the ASID field of the <i>WatchHi</i> register must match the ASID field of the <i>EntryHi</i> register to cause a watch exception.	R/W	Undefined	Required
ASID	2316	ASID value which is required to match that in the <i>EntryHi</i> register if the G bit is zero in the <i>WatchHi</i> register.	R/W	Undefined	Required

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

Fields			Read/		
Name Bits		Description	Write	Reset State	Compliance
		Optional bit mask that qualifies the address in the <i>WatchLo</i> register. If this field is implemented, any bit in this field that is a one inhibits the corresponding address bit from participating in the address match.			
Mask	113	If this field is not implemented, writes to it must be ignored, and reads must return zero.	R/W	Undefined	Optional
		Software may determine how many mask bits are implemented by writing ones the this field and then reading back the result.			
I	2 This bit is set by hardware when an instruction fetch condition matches the values in this watch register pair. When set, the bit remains set until cleared by software, which is accomplished by writing a 1 to the bit.		W1C	Undefined	Required (Release 2)
R 1 I		This bit is set by hardware when a load condition matches the values in this watch register pair. When set, the bit remains set until cleared by software, which is accomplished by writing a 1 to the bit.	W1C	Undefined	Required (Release 2)
W0This bit is set by hardware when a store condition matches the values in this watch register pair. When set, the bit remains set until cleared by software, which is accomplished by writing a 1 to the bit.		W1C	Undefined	Required (Release 2)	
0	2924, 1512	Must be written as zero; returns zero on read.	0	0	Reserved

### Table 8-36 WatchHi Register Field Descriptions

## 8.33 Reserved for Implementations (CP0 Register 22, all Select values)

**Compliance Level:** *Optional: Implementation Dependent.* 

CP0 register 22 is reserved for implementation dependent use and is not defined by the architecture.

## 8.34 Debug Register (CP0 Register 23)

#### **Compliance Level:** *Optional.*

The *Debug* register is part of the EJTAG specification. Refer to that specification for the format and description of this register.

### 8.35 DEPC Register (CP0 Register 24)

#### Compliance Level: Optional.

The *DEPC* register is part of the EJTAG specification. Refer to that specification for the format and description of this register.

All bits of the DEPC register are significant and must be writable.

#### 8.35.1 Special Handling of the DEPC Register in Processors That Implement the MIPS16e ASE

In processors that implement the MIPS16e ASE, a read of the *DEPC* register (via MFC0) returns the following value in the destination GPR:

```
GPR[rt] \leftarrow RestartPC_{31..1} || ISAMode
```

That is, the upper 31 bits of the restart PC are combined with the ISA Mode bit and written to the GPR.

Similarly, a write to the *DEPC* register (via MTC0) takes the value from the GPR and distributes that value to the restart PC and the *ISA Mode* bit, as follows

```
\begin{aligned} \text{RestartPC} &\leftarrow \text{GPR[rt]}_{31..1} \mid \mid \text{ 0} \\ \text{ISAMode} &\leftarrow \text{GPR[rt]}_{0} \end{aligned}
```

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the restart PC, and the lower bit of the restart PC is cleared. The *ISA Mode* bit is loaded from the lower bit of the GPR.

### 8.36 Performance Counter Register (CP0 Register 25)

#### Compliance Level: Recommended.

The MIPS32 Architecture supports implementation dependent performance counters that provide the capability to count events or cycles for use in performance analysis. If performance counters are implemented, each performance counter consists of a pair of registers: a 32-bit control register and a 32-bit counter register. To provide additional capability, multiple performance counters may be implemented.

Performance counters can be configured to count implementation dependent events or cycles under a specified set of conditions that are determined by the control register for the performance counter. The counter register increments once for each enabled event. When the most significant bit of the counter register is a one (the counter overflows), the performance counter optionally requests an interrupt. In implementations of Release 1 of the Architecture, this interrupt is combined in a implementation-dependent way with hardware interrupt 5. In Release 2 of the Architecture, pending interrupts from all performance counters are ORed together to become the PCI bit in the Cause register, and are prioritized as appropriate to the interrupt mode of the processor. Counting continues after a counter register overflow whether or not an interrupt is requested or taken.

Each performance counter is mapped into even-odd select values of the *PerfCnt* register: Even selects access the control register and odd selects access the counter register. Table 8-37 shows an example of two performance counters and how they map into the select values of the *PerfCnt* register.

PerfCnt Performance Counter Value		PerfCnt Register Usage
0	PerfCnt, Select 0	Control Register 0
0	PerfCnt, Select 1	Counter Register 0
1	PerfCnt, Select 2	Control Register 1
1	PerfCnt, Select 3	Counter Register 1

#### Table 8-37 Example Performance Counter Usage of the PerfCnt CP0 Register

More or less than two performance counters are also possible, extending the select field in the obvious way to obtain the desired number of performance counters. Software may determine if at least one pair of Performance Counter Control and Counter registers is implemented via the PC bit in the Config1 register. If the M bit is one in the Performance Counter Counter Control register referenced via a select field of 'n', another pair of Performance Counter Control and Counter registers is implemented at the select values of 'n+2' and 'n+3'.

The Control Register associated with each performance counter controls the behavior of the performance counter. Figure 8-30 shows the format of the Performance Counter Control Register; Table 8-38 describes the Performance Counter Control Register fields.

#### Figure 8-30 Performance Counter Control Register Format

31	30	9 11	10 5	4	3	2	1	0
Μ	W	0	Event	IE	U	S	K	EXL

Fields			Dood/		
Name	Bits	Description	Write	Reset State	Compliance
М	31	If this bit is a one, another pair of Performance Counter Control and Counter registers is implemented at a MTC0 or MFC0 select field value of $(n+2)$ and $(n+3)$ .	R	Preset	Required
W	30	Denotes that the corresponding Counter register is 64 bits wide on a MIPS64 processor. Unused on a MIPS32 processor.	R	Preset	Required
0	2911	Must be written as zero; returns zero on read	0	0	Reserved
Event     105     Selects the even Counter Regiss dependent, but instructions, registructions, counters allow cache miss and events in two cache miss and		Selects the event to be counted by the corresponding Counter Register. The list of events is implementation dependent, but typical events include cycles, instructions, memory reference instructions, branch instructions, cache and TLB misses, etc. Implementations that support multiple performance counters allow ratios of events, e.g., cache miss ratios if cache miss and memory references are selected as the events in two counters	R/W	Undefined	Required
IE	4	Interrupt Enable. Enables the interrupt request when the corresponding counter overflows (the most significant bit of the counter is one. This is bit 31 for a 32-bit wide counter or bit 63 of a 64-bit wide counter, denoted by the W bit in this register).Note that this bit simply enables the interrupt request. The actual interrupt is still gated by the normal interrupt masks and enable in the <i>Status</i> register.EncodingMeaning 0 Performance counter interrupt disabled 11Performance counter interrupt enabled	R/W	0	Required
U	3	Enables event counting in User Mode. Refer to Section         Section 3.4, "User Mode" on page 10 for the conditions         under which the processor is operating in User Mode.         Encoding       Meaning         0       Disable event counting in User Mode         1       Enable event counting in User Mode	R/W	Undefined	Required
S	2	Enables event counting in Supervisor Mode (for those processors that implement Supervisor Mode). Refer to Section Section 3.3, "Supervisor Mode" on page 9 for the conditions under which the processor is operating 	R/W	Undefined	Required

### Table 8-38 Performance Counter Control Register Field Descriptions

Fields				Read/		
Name	Bits		Description	Write	Reset State	Compliance
K	1	Enables even usual defini Section 3.2 event count <i>Status</i> regis	ent counting in Kernel Mode. Unlike the tion of Kernel Mode as described in Section , "Kernel Mode" on page 9, this bit enables ing only when the EXL and ERL bits in the ter are zero.	R/W	Undefined	Required
		Encoding	Meaning			
		0	Disable event counting in Kernel Mode			
		1	Enable event counting in Kernel Mode			
		Enables event counting when the EXL bit in the <i>Status</i> register is one and the ERL bit in the <i>Status</i> register is zero.				
		Encoding	Meaning		Undefined	
EXL	0	0	Disable event counting while $EXL = 1$ , ERL = 0	R/W		Required
		1	Enable event counting while $EXL = 1$ , ERL = 0			-
		Counting is Status regis	s never enabled when the ERL bit in the ter or the DM bit in the <i>Debug</i> register is			

#### Table 8-38 Performance Counter Control Register Field Descriptions

The Counter Register associated with each performance counter increments once for each enabled event. Figure 8-31 shows the format of the Performance Counter Counter Register; Table 8-39 describes the Performance Counter Counter Register fields.

#### Figure 8-31 Performance Counter Counter Register Format

31	(	)
	Event Count	

#### Table 8-39 Performance Counter Counter Register Field Descriptions

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
Event Count	310	Increments once for each event that is enabled by the corresponding Control Register. When the most significant bit is one, a pending interrupt request is ORed with those from other performance counters and indicated by the PCI bit in the <i>Cause</i> register.	R/W	Undefined	Required

### 8.37 ErrCtl Register (CP0 Register 26, Select 0)

#### **Compliance Level:** Optional.

The *ErrCtl* register provides an implementation dependent diagnostic interface with the error detection mechanisms implemented by the processor. This register has been used in previous implementations to read and write parity or ECC information to and from the primary or secondary cache data arrays in conjunction with specific encodings of the Cache instruction or other implementation-dependent method. The exact format of the ErrCtl register is implementation dependent and not specified by the architecture. Refer to the processor specification for the format of this register and a description of the fields.

### 8.38 CacheErr Register (CP0 Register 27, Select 0)

#### **Compliance Level:** *Optional.*

The CacheErr register provides an interface with the cache error detection logic that may be implemented by a processor.

The exact format of the *CacheErr* register is implementation dependent and not specified by the architecture. Refer to the processor specification for the format of this register and a description of the fields.

### 8.39 TagLo Register (CP0 Register 28, Select 0, 2)

Compliance Level: Required if a cache is implemented; Optional otherwise.

The *TagLo* and *TagHi* registers are read/write registers that act as the interface to the cache tag array. The Index Store Tag and Index Load Tag operations of the CACHE instruction use the *TagLo* and *TagHi* registers as the source or sink of tag information, respectively.

The exact format of the *TagLo* and *TagHi* registers is implementation dependent. Refer to the processor specification for the format of this register and a description of the fields.

However, software must be able to write zeros into the *TagLo* and *TagHi* registers and then use the Index Store Tag cache operation to initialize the cache tags to a valid state at powerup.

It is implementation dependent whether there is a single *TagLo* register that acts as the interface to all caches, or a dedicated *TagLo* register for each cache. If multiple *TagLo* registers are implemented, they occupy the even select values for this register encoding, with select 0 addressing the instruction cache and select 2 addressing the data cache. Whether individual *TagLo* registers are implemented or not for each cache, processors must accept a write of zero to select 0 and select 2 of *TagLo* as part of the software process of initializing the cache tags at powerup.

### 8.40 DataLo Register (CP0 Register 28, Select 1, 3)

#### **Compliance Level:** Optional.

The *DataLo* and *DataHi* registers are read-only registers that act as the interface to the cache data array and are intended for diagnostic operation only. The Index Load Tag operation of the CACHE instruction reads the corresponding data values into the *DataLo* and *DataHi* registers.

The exact format and operation of the *DataLo* and *DataHi* registers is implementation dependent. Refer to the processor specification for the format of this register and a description of the fields.

It is implementation dependent whether there is a single *DataLo* register that acts as the interface to all caches, or a dedicated *DataLo* register for each cache. If multiple *DataLo* registers are implemented, they occupy the odd select values for this register encoding, with select 1 addressing the instruction cache and select 3 addressing the data cache.

### 8.41 TagHi Register (CP0 Register 29, Select 0, 2)

Compliance Level: Required if a cache is implemented; Optional otherwise.

The *TagLo* and *TagHi* registers are read/write registers that act as the interface to the cache tag array. The Index Store Tag and Index Load Tag operations of the CACHE instruction use the *TagLo* and *TagHi* registers as the source or sink of tag information, respectively.

The exact format of the *TagLo* and *TagHi* registers is implementation dependent. Refer to the processor specification for the format of this register and a description of the fields. However, software must be able to write zeros into the *TagLo* and *TagHi* registers and the use the Index Store Tag cache operation to initialize the cache tags to a valid state at powerup.

It is implementation dependent whether there is a single *TagHi* register that acts as the interface to all caches, or a dedicated *TagHi* register for each cache. If multiple *TagHi* registers are implemented, they occupy the even select values for this register encoding, with select 0 addressing the instruction cache and select 2 addressing the data cache. Whether individual *TagHi* registers are implemented or not for each cache, processors must accept a write of zero to select 0 and select 2 of *TagHi* as part of the software process of initializing the cache tags at powerup.

### 8.42 DataHi Register (CP0 Register 29, Select 1, 3)

#### **Compliance Level:** *Optional.*

The *DataLo* and *DataHi* registers are read-only registers that act as the interface to the cache data array and are intended for diagnostic operation only. The Index Load Tag operation of the CACHE instruction reads the corresponding data values into the *DataLo* and *DataHi* registers.

The exact format and operation of the *DataLo* and *DataHi* registers is implementation dependent. Refer to the processor specification for the format of this register and a description of the fields.

### 8.43 ErrorEPC (CP0 Register 30, Select 0)

#### Compliance Level: Required.

The *ErrorEPC* register is a read-write register, similar to the *EPC* register, except that *ErrorEPC* is used on error exceptions. All bits of the *ErrorEPC* register are significant and must be writable. It is also used to store the program counter on Reset, Soft Reset, Nonmaskable Interrupt (NMI), and Cache Error exceptions.

The *ErrorEPC* register contains the virtual address at which instruction processing can resume after servicing an error. *ErrorEPC* contains either:

- the virtual address of the instruction that was the direct cause of the exception, or
- the virtual address of the immediately preceding branch or jump instruction when the error causing instruction is in a branch delay slot.

Unlike the EPC register, there is no corresponding branch delay slot indication for the ErrorEPC register.

Figure 8-32 shows the format of the *ErrorEPC* register; Table 8-40 describes the *ErrorEPC* register fields.

#### Figure 8-32 ErrorEPC Register Format

31	0
ErrorEPC	

#### **Table 8-40 ErrorEPC Register Field Descriptions**

Fields			Read/		
Name	Bits	Description	Write	Reset State	Compliance
ErrorEPC	310	Error Exception Program Counter	R/W	Undefined	Required

#### 8.43.1 Special Handling of the ErrorEPC Register in Processors That Implement the MIPS16e ASE

In processors that implement the MIPS16e ASE, a read of the *ErrorEPC* register (via MFC0) returns the following value in the destination GPR:

 $GPR[rt] \leftarrow RestartPC_{31..1} || ISAMode$ 

That is, the upper 31 bits of the restart PC are combined with the ISA Mode bit and written to the GPR.

Similarly, a write to the *ErrorEPC* register (via MTC0) takes the value from the GPR and distributes that value to the restart PC and the *ISA Mode* bit, as follows

 $\begin{aligned} \text{RestartPC} &\leftarrow \text{GPR[rt]}_{31..1} || & 0\\ \text{ISAMode} &\leftarrow \text{GPR[rt]}_{0} \end{aligned}$ 

That is, the upper 31 bits of the GPR are written to the upper 31 bits of the restart PC, and the lower bit of the restart PC is cleared. The *ISA Mode* bit is loaded from the lower bit of the GPR.

### 8.44 DESAVE Register (CP0 Register 31)

#### **Compliance Level:** *Optional.*

The *DESAVE* register is part of the EJTAG specification. Refer to that specification for the format and description of this register.

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

# Alternative MMU Organizations

The main body of this specification describes the TLB-based MMU organization. This appendix describes other potential MMU organizations.

### A.1 Fixed Mapping MMU

As an alternative to the full TLB-based MMU, the MIPS32 Architecture supports a lightweight memory management mechanism with fixed virtual-to-physical address translation, and no memory protection beyond what is provided by the address error checks required of all MMUs. This may be useful for those applications which do not require the capabilities of a full TLB-based MMU.

#### A.1.1 Fixed Address Translation

Address translation using the Fixed Mapping MMU is done as follows:

- Kseg0 and Kseg1 addresses are translated in an identical manner to the TLB-based MMU: they both map to the low 512MB of physical memory.
- Useg/Suseg/Kuseg addresses are mapped by adding 1GB to the virtual address when the ERL bit is zero in the Status register, and are mapped using an identity mapping when the ERL bit is one in the Status register.
- Sseg/Ksseg/Kseg2/Kseg3 addresses are mapped using an identity mapping.

Supervisor Mode is not supported with a Fixed Mapping MMU.

Table 8-41 lists all mappings from virtual to physical addresses. Note that address error checking is still done before the translation process. Therefore, an attempt to reference kseg0 from User Mode still results in an address error exception, just as it does with a TLB-based MMU.

Segment		Generates Physical Address			
Name	Virtual Address	$Status_{ERL} = 0$	Status <sub>ERL</sub> = 1		
useg	16#0000 0000	16#4000 0000	16#0000 0000		
suseg	through	through	through		
kuseg	16#7FFF FFFF	16#BFFF FFFF	16#7FFF FFFF		
	16#8000 0000	16#0000 0000			
kseg0	through	through			
	16#9FFF FFFF	16#1FF	F FFFF		
	16#A000 0000	16#000	0 0000		
lrang 1	through	thro	ugh		
KSCg1	16#BFFF FFFF	16#16#1FFF FFFF			

**Table 8-41 Physical Address Generation from Virtual Addresses** 

Segment		Generates Physical Address				
Name Virtual Address		Status <sub>ERL</sub> = 0	Status <sub>ERL</sub> = 1			
sseg	16#C000 0000	16#C000 0000				
ksseg	through	through				
kseg2	16#DFFF FFFF	16#DFFF FFFF				
	16#E000 0000	16#E00	0 0000			
kseg3	through	thro	ough			
	16#FFFF FFFF	16#FFF	F FFFF			

Note that this mapping means that physical addresses 16#2000 0000 through 16#3FFF FFFF are inaccessible when the ERL bit is off in the *Status* register, and physical addresses 16#8000 0000 through 16#BFFF FFFF are inaccessible when the ERL bit is on in the *Status* register.

Figure 8-33 shows the memory mapping when the ERL bit in the *Status* register is zero; Figure 8-34 shows the memory mapping when the ERL bit is one.





#### A.1.2 Cacheability Attributes

Because the TLB provided the cacheability attributes for the kuseg, kseg2, and kseg3 segments, some mechanism is required to replace this capability when the fixed mapping MMU is used. Two additional fields are added to the *Config* register whose encoding is identical to that of the K0 field. These additions are the K23 and KU fields which control the cacheability of the kseg2/kseg3 and the kuseg segments, respectively. Note that when the ERL bit is on in the *Status* register, kuseg data references are always treated as uncacheable references, independent of the value of the KU field. The operation of the processor is **UNDEFINED** if the ERL bit is set while the processor is executing instructions from kuseg.

The cacheability attributes for kseg0 and kseg1 are provided in the same manner as for a TLB-based MMU: the cacheability attribute for kseg0 comes from the K0 field of *Config*, and references to kseg1 are always uncached.

Figure 8-35 shows the format of the additions to the Config register; Table 8-42 describes the new Config register fields.

	Figure 8-35 Config Register Additions										
31	30 28	27 25	24 16	15	14 13	12 10	9 7	6	3	2	0
Μ	K23	KU	0	BE	AT	AR	MT	0		K0	

#### **Table 8-42 Config Register Field Descriptions**

Fields			Read/		
Name	Bits	Description Write Res		Reset State	Compliance
K23	30:28	Kseg2/Kseg3 coherency algorithm. See Table 8-8 on page 61 for the encoding of this field.	R/W	Undefined	Optional
KU	27:25	Kuseg coherency algorithm when $Status_{ERL}$ is zero. See Table 8-8 on page 61 for the encoding of this field.	R/W	Undefined	Optional

#### A.1.3 Changes to the CP0 Register Interface

Relative to the TLB-based address translation mechanism, the following changes are necessary to the CP0 register interface:

- The Index, Random, EntryLo0, EntryLo1, Context, PageMask, Wired, and EntryHi registers are no longer required and may be removed.
- The TLBWR, TLBWI, TLBP, and TLBR instructions are no longer required and should cause a Reserved Instruction Exception.

### A.2 Block Address Translation

This section describes the architecture for a block address translation (BAT) mechanism that reuses much of the hardware and software interface that exists for a TLB-Based virtual address translation mechanism. This mechanism has the following features:

- It preserves as much as possible of the TLB-Based interface, both in hardware and software.
- It provides independent base-and-bounds checking and relocation for instruction references and data references.
- It provides optional support for base-and-bounds relocation of kseg2 and kseg3 virtual address regions.

### A.2.1 BAT Organization

The BAT is an indexed structure which is used to translate virtual addresses. It contains pairs of instruction/data entries which provide the base-and-bounds checking and relocation for instruction references and data references, respectively. Each entry contains a page-aligned bounds virtual page number, a base page frame number (whose width is implementation dependent), a cache coherence field (C), a dirty (D) bit, and a valid (V) bit. Figure 8-36 shows the logical arrangement of a BAT entry.



The BAT is indexed by the reference type and the address region to be checked as shown in Table 8-43.

Entry Index	Reference Type	Address Region
0	Instruction	usag/kusag
1	Data	useg/kuseg
2	Instruction	kseg2
3	Data	(or kseg2 and kseg3)
4	Instruction	ksag3
5	Data	KSeg3

Table	8-43	BAT	Entry	Assignments
-------	------	-----	-------	-------------

Entries 0 and 1 are required. Entries 2, 3, 4 and 5 are optional and may be implemented as necessary to address the needs of the particular implementation. If entries for kseg2 and kseg3 are not implemented, it is implementation-dependent how, if at all, these address regions are translated. One alternative is to combine the mapping for kseg2 and kseg3 into a single pair of instruction/data entries. Software may determine how many BAT entries are implemented by looking at the MMU Size field of the *Config1* register.

### A.2.2 Address Translation

When a virtual address translation is requested, the BAT entry that is appropriate to the reference type and address region is read. If the virtual address is greater than the selected bounds address, or if the valid bit is off in the entry, a TLB Invalid exception of the appropriate reference type is initiated. If the reference is a store and the D bit is off in the entry, a TLB Modified exception is initiated. Otherwise, the base PFN from the selected entry, shifted to align with bit 12, is added to the virtual address to form the physical address. The BAT process can be described as follows:

```
\begin{split} \label{eq:approx_eq} i &\leftarrow \text{SelectIndex (reftype, va)} \\ \text{bounds} &\leftarrow \text{BAT[i]}_{\text{BoundsVPN}} \mid \mid 1^{12} \\ \text{pfn} &\leftarrow \text{BAT[i]}_{\text{BasePFN}} \\ \text{c} &\leftarrow \text{BAT[i]}_{\text{C}} \\ \text{d} &\leftarrow \text{BAT[i]}_{\text{D}} \\ \text{v} &\leftarrow \text{BAT[i]}_{\text{V}} \\ \text{if (va > bounds) or (v = 0) then} \\ &\quad \text{InitiateTLBInvalidException(reftype)} \\ \text{endif} \\ \text{if (d = 0) and (reftype = store) then} \\ &\quad \text{InitiateTLBModifiedException()} \\ \text{endif} \end{split}
```

 $pa \leftarrow va + (pfn || 0^{12})$ 

Making all addresses out-of-bounds can only be done by clearing the valid bit in the BAT entry. Setting the bounds value to zero leaves the first virtual page mapped.

#### A.2.3 Changes to the CP0 Register Interface

Relative to the TLB-based address translation mechanism, the following changes are necessary to the CP0 register interface:

- The Index register is used to index the BAT entry to be read or written by the TLBWI and TLBR instructions.
- The EntryHi register is the interface to the BoundsVPN field in the BAT entry.
- The *EntryLo0* register is the interface to the BasePFN and C, D, and V fields of the BAT entry. The register has the same format as for a TLB-based MMU.
- The *Random*, *EntryLo1*, *Context*, *PageMask*, and *Wired* registers are eliminated. The effects of a read or write to these registers is **UNDEFINED**.
- The TLBP and TLBWR instructions are unnecessary. The TLBWI and TLBR instructions reference the BAT entry whose index is contained in the *Index* register. The effects of executing a TLBP or TLBWR are **UNDEFINED**, but processors should prefer a Reserved Instruction Exception.

# **Revision History**

In the left hand page margins of this document you may find vertical change bars to note the location of significant changes to this document since its last release. Significant changes are defined as those which you should take note of as you use the MIPS IP. Changes to correct grammar, spelling errors or similar may or may not be noted with change bars. Change bars will be removed for changes which are more than one revision old.

Please note: Limitations on the authoring tools make it difficult to place change bars on changes to figures. Change bars on figure titles are used to denote a potential change in the figure itself.

Revision Date		Description		
0.92	January 20, 2001	Internal review copy of reorganized and updated architecture documentati		
0.95	March 12, 2001	Clean up document for external review release		
		Update based on review feedback:		
		Change ProbEn to ProbeTrap in the EJTAG Debug entry vector location discussion.		
		• Add cache error and EJTAG Debug exceptions to the list of exceptions that do not go through the general exception processing mechanism.		
		• Fix incorrect branch offset adjustment in general exception processing pseudo code to deal with extended MIPS16e instructions.		
		- Add $\mathrm{Config}_{\mathrm{VI}}$ to denote an instruction cache with both virtual indexing and virtual tags.		
1.00	August 29, 2002	• Correct XContext register description to note that both BadVPN2 and R fields are UNPREDICTABLE after an address error exception.		
		• Note that Supervisor Mode is not supported with a Fixed Mapping MMU.		
		• Define TagLo bits 43 as implementation dependent.		
		• Describe the intended usage model differences between Reset and Soft Reset Exceptions.		
		• Correct the minimum number of TLB entries to be 3, not 2, and show an example of the need for 3.		
		<ul> <li>Modify the description of PageMask and the TLB lookup process to acknowledge the fact that not all implementations may support all page sizes.</li> </ul>		
		Update the specification with the changes introduced in Release 2 of the Architecture. Changes in this revision include:		
	September 1, 2002	• The following new Coprocessor 0 registers were added: EBase, HWREna, IntCtl, PageGrain, SRSCtl, SRSMap.		
1.90		<ul> <li>The following Coprocessor 0 registers were modified: Cause, Config, Config2, Config3, EntryHi, EntryLo0, EntryLo1, PageMask, PerfCnt, Status, WatchHi, WatchLo.</li> </ul>		
		• The descriptions of Virtual memory, exceptions, and hazards have been updated to reflect the changes in Release 2.		
		• A chapter on GPR shadow regsiters has been added.		
		• The chapter on CP0 hazards has been completely rewriten to reflect the Release 2 changes.		

MIPS32<sup>™</sup> Architecture For Programmers Volume III, Revision 2.00

Revision	Date	Description		
		Complete the update to include Release 2 changes. These include:		
		• Make bits 1211 of the PageMask register power up zero and be gated by 1K page enable. This eliminates the problem of having these bits set to 2#11 on a Release 2 chip in which kernel software has not enabled 1K page support.		
2.00		• Correct the address of the cache error vector when the BEV bit is 1. It should be 16#BFC0.0300,. not 16#BFC0.0200.		
		• Correct the introduction to shadow registers to note that the SRSCtl register is not updated at the end of an exception in which $\text{Status}_{\text{BEV}} = 1$ .		
	June 9, 2003	• Clarify that a MIPS16e PC-relative load reference is a data reference for the purposes of the Watch registers.		
		• Add note about a hardware interrupt being deasserted between the time that the processor detects the interrupt request and the time that the software interrupt handler runs. Software must be prepared for this case and simply dismiss the interrupt via an ERET.		
		• Add restriction that software must set EBase <sub>1512</sub> to zero in all bit positions less than or equal to the most significant bit in the vector offset. This is only required in certain combinations of vector number and vector spacing when using VI or EIC Interrupt modes.		
		• Add suggested software TLB init routine which reduced the probability of triggering a machine check.		