MIPS® Architecture For Programmers Volume III: The MIPS64® and microMIPS64™ Privileged Resource Architecture

Document Number: MD00091
Revision 3.12
April 28, 2011

MIPS Technologies, Inc.
955 East Arques Avenue
Sunnyvale, CA 94085-4521

Template: nB1.03, Built with tags: 2B ARCH FPU_PS FPU_PSandARCH MIPS64

Contents

Figures

Tables

Chapter 1

About This Book

The MIPS® Architecture For Programmers Volume III: The MIPS64® and microMIPS64™ Privileged Resource Architecture comes as part of a multi-volume set.

- Volume I-A describes conventions used throughout the document set, and provides an introduction to the MIPS64® Architecture
- Volume I-B describes conventions used throughout the document set, and provides an introduction to the microMIPS64™ Architecture
- Volume II-A provides detailed descriptions of each instruction in the MIPS64® instruction set
- Volume II-B provides detailed descriptions of each instruction in the microMIPS64™ instruction set
- Volume III describes the MIPS64® and microMIPS64™ Privileged Resource Architecture which defines and governs the behavior of the privileged resources included in a MIPS® processor implementation
- Volume IV-a describes the MIPS16e™ Application-Specific Extension to the MIPS64® Architecture. Beginning with Release 3 of the Architecture, microMIPS is the preferred solution for smaller code size.
- Volume IV-b describes the MDMX™ Application-Specific Extension to the MIPS64® Architecture and microMIPS64™.
- Volume IV-c describes the MIPS-3D® Application-Specific Extension to the MIPS® Architecture
- Volume IV-d describes the SmartMIPS®Application-Specific Extension to the MIPS32® Architecture and the microMIPS32™ Architecture and is not applicable to the MIPS64® document set nor the microMIPS64™ document set
- Volume IV-e describes the MIPS® DSP Application-Specific Extension to the MIPS® Architecture
- Volume IV-f describes the MIPS® MT Application-Specific Extension to the MIPS® Architecture
- Volume IV-h describes the MIPS® MCU Application-Specific Extension to the MIPS® 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

1.2.3 UNSTABLE

UNSTABLE results or values may vary as a function of time on the same implementation or instruction. Unlike **UNPREDICTABLE** values, software may depend on the fact that a sampling of an **UNSTABLE** value results in a legal transient value that was correct at some point in time prior to the sampling.

UNSTABLE values have one implementation restriction:

- Implementations of operations generating **UNSTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode

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.

Table 1.1 Symbols Used in Instruction Operation Statements

Symbol	Meaning
←	Assignment
=, ≠	Tests for equality and inequality
||	Bit string concatenation
x^y	A *y*-bit string formed by *y* copies of the single-bit value *x*
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.
0bn	A constant value *n* in base *2*. For instance 0b100 represents the binary value 100 (decimal 4).
0xn	A constant value *n* in base *16*. For instance 0x100 represents the hexadecimal value 100 (decimal 256).
$x_{y..z}$	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

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
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'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 *GPR[0]* is always zero. In Release 2 of the Architecture, GPR[x] is a short-hand notation for $SGPR[\ SRSCtl_{CSS},\ x]$.
SGPR[s,x]	In Release 2 of the Architecture and subsequent releases, multiple copies of the CPU general-purpose registers may be implemented. *SGPR[s,x]* refers to GPR set *s*, register *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 *x*
CPR[z,x,s]	Coprocessor unit *z*, general register *x*, select *s*
CP2CPR[x]	Coprocessor unit 2, general register *x*
CCR[z,x]	Coprocessor unit *z*, control register *x*
CP2CCR[x]	Coprocessor unit 2, control register *x*
COC[z]	Coprocessor unit *z* condition signal
Xlat[x]	Translation of the MIPS16e GPR number *x* into the corresponding 32-bit GPR number
BigEndianMem	Endian mode as configured at chip reset (0 →Little-Endian, 1 → 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 → Little-Endian, 1 → Big-Endian). In User mode, this endianness may be switched by setting the *RE* bit in the *Status* register. Thus, BigEndianCPU may be computed as (BigEndianMem XOR ReverseEndian).
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 *RE* bit of the *Status* register. Thus, ReverseEndian may be computed as (SR_{RE} and User mode).
LLbit	Bit of **virtual** state used to specify operation for instructions that provide atomic read-modify-write. *LLbit* is set when a linked load occurs and is tested 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.

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
I:, **I+n**:, **I-n**:	This occurs as a prefix to *Operation* 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 **I**. 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 **I**, 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 operation description that writes the result register in a section labeled **I+1**. The effect of pseudocode statements for the current instruction labelled **I+1** appears to occur "at the same time" as the effect of pseudocode statements labeled **I** 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.
PC	The *Program Counter* 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 *PC* during an instruction time. If no value is assigned to *PC* 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 *PC* during the instruction time of the instruction in the branch delay slot. In the MIPS Architecture, the PC value is only visible indirectly, such as when the processor stores the restart address into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception. The PC value contains a full 64-bit address all of which are significant during a memory reference.
ISA Mode	In processors that implement the MIPS16e Application Specific Extension or the microMIPS base architectures, the *ISA Mode* is a single-bit register that determines in which mode the processor is executing, as follows: **Encoding** / **Meaning** 0 — The processor is executing 32-bit MIPS instructions 1 — The processor is executing MIIPS16e instructions In the MIPS Architecture, the ISA Mode value is only visible indirectly, such as when the processor stores a combined value of the upper bits of PC and the ISA Mode into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception.
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.
SEGBITS	The number of virtual address bits implemented in a segment of the address space is represented by the symbol SEGBITS. As such, if 40 virtual address bits are implemented in a segment, the size of the segment is $2^{SEGBITS} = 2^{40}$ bytes.
FP32RegistersMode	Indicates whether the FPU has 32-bit or 64-bit floating point registers (FPRs). In MIPS32 Release 1, the FPU has 32 32-bit FPRs in which 64-bit data types are stored in even-odd pairs of FPRs. In MIPS64, (and optionally in MIPS32 Release2 and MIPSr3) the FPU has 32 64-bit FPRs in which 64-bit data types are stored in any FPR. In MIPS32 Release 1 implementations, **FP32RegistersMode** 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 **FP32RegisterMode** is computed from the FR bit in the *Status* 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 **FP32RegistersMode** is computed from the FR bit in the *Status* register.

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
InstructionInBranchDelaySlot	Indicates whether the instruction at the Program Counter address was executed in the delay slot of a branch or jump. This condition reflects the *dynamic* state of the instruction, not the *static* 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(exception, 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.

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

For comments or questions on the MIPS64® Architecture or this document, send Email to support@mips.com.

Chapter 2

The MIPS64 and microMIPS64 Privileged Resource Architecture

2.1 Introduction

The MIPS64 and microMIPS64 Privileged Resource Architecture (PRA) is a set of environments and capabilities on which the Instruction Set Architectures operate. 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 9.

Chapter 3

MIPS64 and microMIPS64 Operating Modes

The MIPS64 and microMIPS64 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 MIPS 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 MIPS64 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. Release 3 (e.g. MIPSr3) introduced the microMIPS instruction set, so all microMIPS processors may implement a 64-bit floating point unit.

Finally, the MIPS64 and microMIPS64 PRA provides backward compatible support for 32-bit programs by providing enables for both 64-bit addressing and 64-bit operations. If access is not enabled, an attempt to reference a 64-bit address or an instruction that implements a 64-bit operation results in an exception.

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 0b00
- 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 0b01
- 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 0b10
- The EXL and ERL bits in the *Status* register are both zero

3.5 Other Modes

3.5.1 64-bit Address Enable

Access to 64-bit addresses are enabled under any of the following conditions:

- A legal reference to a kernel address space occurs and the KX bit in the *Status* register is a one
- A legal reference to a supervisor address space occurs and the SX bit in the *Status* register is a one
- A legal reference to a user address space occurs and the UX bit in the *Status* register is a one

Note that the operating mode of the processor is not relevant to 64-bit address enables. That is, a reference to user address space made while the processor is operating in Kernel Mode is controlled by the state of the UX bit, not by the KX bit.

An attempt to reference a 64-bit address space when 64-bit addresses are not enabled results in an Address Error Exception (either AdEL or AdES, depending on the type of reference).

When a TLB miss occurs, the choice of the Exception Vector is also determined by the 64-bit address enable[1]. If 64-bit addresses are not enabled for the reference, the TLB Refill Vector is used. If 64-bit addresses are enabled for the reference, the XTLB Refill Vector is used.

3.5.2 64-bit Operations Enable

Instructions that perform 64-bit operations are legal under any of the following conditions:

1 For ksseg/sseg access while in supervisor mode, please refer to Note 2 of Table 4.2.

- The processor is operating in Kernel Mode, Supervisor Mode, or Debug Mode, as described above.
- The PX bit in the *Status* register is a one
- The processor is operating in User Mode, as described above, and the UX bit in the *Status* register is a one.

The last two bullets imply that 64-bit operations are legal in User Mode when either the PX bit or the UX bit is a one in the *Status* register.

An attempt to execute an instruction which performs 64-bit operations when such instructions are not enabled results in a Reserved Instruction Exception.

3.5.3 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 enabled only if 64-bit operations enabled.
- In an implementation of Release 2 (and subsequent releases) 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. Release 3 (e.g. MIPSr3) introduced the microMIPS instruction set. So on all microMIPS processors, 64-bit floating point operations are enabled if the F64 bit in the *FIR* register is a one .

3.5.4 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. 64-bit FPRs are supported for all processors using Architecture releases subsequent to Release 2, including all microMIPS processors.

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

3.5.5 Coprocessor 0 Enable

Access to Coprocessor 0 registers are enabled under any of the following conditions:

- The processor is running in Kernel Mode or Debug Mode, as defined above
- The CU0 bit in the *Status* register is one.

3.5.6 ISA Mode

Release 3 of the Architecture (e.g. MIPSr3™) introduced a second branch of the instruction set family, microMIPS64. Devices can implement both ISA branches (MIPS64 and microMIPS64) or only one branch.

The ISA Mode bit is used to denote which ISA branch to use when decoding instructions. This bit is normally not visible to software. It's value is saved to any GPR that would be used as a jump target address, such as GPR31 when written by a JAL instruction or the source register for a JR instruction.

For processors that implement the MIPS64 ISA, the ISA Mode bit value of zero selects MIPS64. For processors that implement the microMIPS64 ISA, the ISA Mode bit value of one selects microMIPS64. For processors that implement the MIPS16e™ ASE, the ISA Mode bit value of one selects MIPS16e. A processor is not allowed to implement both MIPS16e and microMIPS.

Please read Volume II-B: Introduction to the microMIPS64 Instruction Set, Section 5.3, "ISA Mode Switch" for a more in-depth description of ISA mode switching between the ISA branches and the ISA Mode bit.

Chapter 4

Virtual Memory

4.1 Differences between Releases 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 (and subsequent releases), 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™ 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 $Config3_{SP}$ bit and enabled by the $PageGrain_{ESP}$ bit.

4.1.2 Physical Memory

In Release 1 of the Architecture, the physical address size was limited by the format of the *EntryLo0* and *EntryLo1* registers to 36 bits. Some applications of MIPS processors already require more than 36 bits of physical address (for example, high-end networking), and others are expected to appear during the lifetime of Release 2 of the architecture. As such, Release 2 added an optional extension to the architecture to provide up to 59 bits of physical address for MIPS64 processors. This extension is optional because several operating systems currently use the reserved bits to the left of the PFN field in the *EntryLo0* and *EntryLo1* registers for PTE software flags. The flags are loaded directly into these registers on a TLB Refill exception. As such, for compatibility with existing software, the extension of the PFN field must be done with an explicit enable.

Support for extended PFNs is denoted by the $Config3_{LPA}$ bit and enabled by the $PageGrain_{ELPA}$ bit.

4.1.3 Protection of Virtual Memory Pages

In Release 3 of the Architecture, e.g. MIPSr3, two optional control bits are added to each TLB entry. These bits, *RI (Read Inhibit) and XI (Execute Inhibit)*, allows more types of protection to be used for virtual pages - including write-only pages, non-executable pages.

This feature originated in the SmartMIPS ASE but has been modified from the original SmartMIPS definition. For the Release 3 version of this feature, each of the RI and XI bits can be separately implemented. For the Release 3 version of this feature, new exception codes are used when a TLB access does not obey the RI/XI bits.

4.1.4 Context Register

In Release 3 of the Architecture, e.g. MIPSr3, the *Context/XContext* registers are a read/write registers containing a address pointer that can point to an arbitrary power-of-two aligned data structure in memory, such as an entry in the page table entry (PTE) array. In Releases 1 & 2, this pointer was defined to reference a fixed-sized 16-byte structure in memory within a linear array containing an entry for each even/odd virtual page pair. The Release 3 version of the *Context/XContext* registers can be used far more generally.

This feature originated in the SmartMIPS ASE. This feature is optional in the Release 3 version of the base architecture.

4.2 Terminology

4.2.1 Address Space

An *Address Space* is the range of all possible addresses that can be generated for a particular addressing mode. There is one 64-bit Address Space and one 32-bit Compatibility Address Space that is mapped into a subset of the 64-bit Address Space.

4.2.2 Segment and Segment Size (SEGBITS)

A *Segment* is a defined subset of an Address Space that has self-consistent reference and access behavior. A 32-bit Compatibility Segment is part of the 32-bit Compatibility Address Space and is either 2^{29} or 2^{31} bytes in size, depending on the specific Segment. A 64-bit Segment is part of the 64-bit Address Space and is no larger than 2^{62} bytes in size, but may be smaller on an implementation dependent basis. The symbol *SEGBITS* is used to represent the actual number of bits implemented in each 64-bit Segment. As such, if 40 virtual address bits were implemented, the actual size of the Segment would be $2^{SEGBITS} = 2^{40}$ bytes. Software may determine the value of SEGBITS by writing all ones to the *EntryHi* register and reading the value back. Bits read as "1" from the VPN2 field allow software to determine the boundary between the VPN2 and Fill fields to calculate the value of SEGBITS.

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 Fill fields to calculate the value of PABITS.

4.3 Virtual Address Spaces

With support for 64-bit operations and address calculation, the MIPS64/microMIPS64 architecture implicitly defines and provides support for a 64-bit virtual Address Space, sub-divided into four Segments selected by bits 63..62 of the virtual address. To provide compatibility for 32-bit programs and MIPS32/microMIPS32 processors, a 2^{32}-byte Compatibility Address Space is defined, separated into two non-contiguous ranges in which the upper 32 bits of the 64-bit

address are the sign extension of bit 31. The Compatibility Address Space is similarly sub-divided into Segments selected by bits 31..29 of the virtual address. Figure 4-1 shows the layout of the Address Spaces, including the Compatibility Address Space and the segmentation of each Address Space.

Figure 4-1 Virtual Address Space

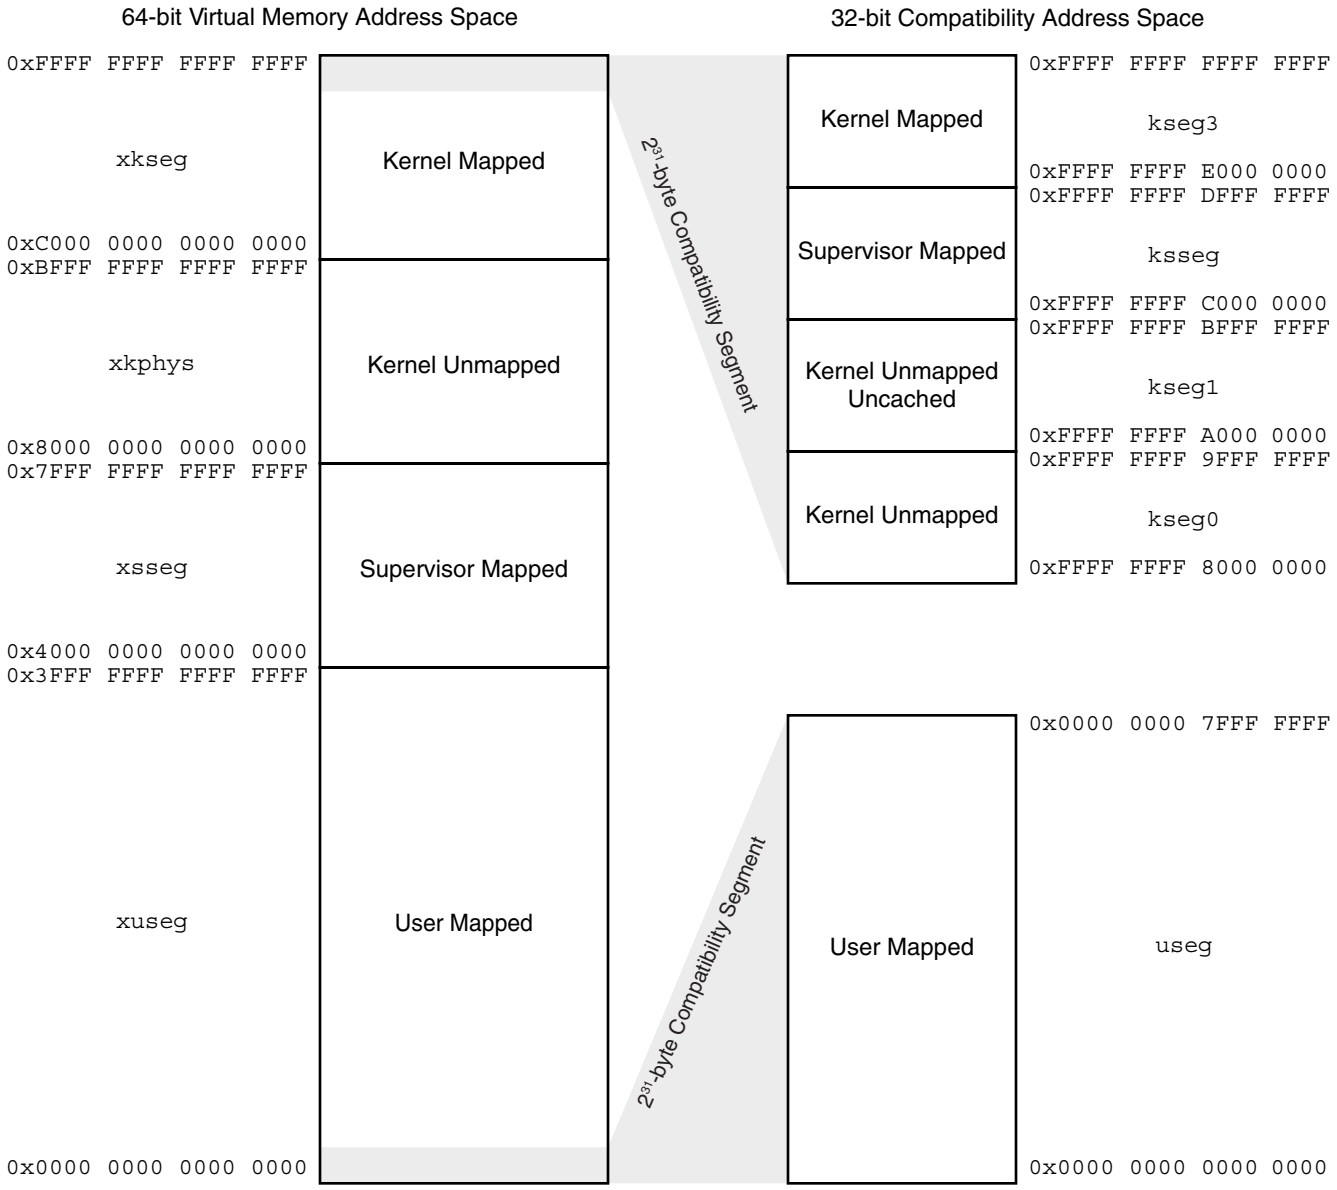

Each Segment of an Address Space is classified as “Mapped” or “Unmapped”. A “Mapped” address is one that is translated through the TLB or other address translation unit. An “Unmapped” address is one which is not translated through the TLB and which provides a window into the lowest portion of the physical 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.

Table 4.1 Virtual Memory Address Spaces

$VA_{63..62}$	Segment Name(s)	Maximum Address Range	Associated with Mode	Reference Legal from Mode(s)	Actual Segment Size	64-bit Address Enable	Segment Type
0b11	kseg3	0xFFFF FFFF FFFF FFFF through 0xFFFF FFFF E000 0000	Kernel	Kernel	2^{29} bytes	Always	32-bit Compatibility
	sseg ksseg	0xFFFF FFFF DFFF FFFF through 0xFFFF FFFF C000 0000	Supervisor	Supervisor Kernel	2^{29} bytes	Always	32-bit Compatibility
	kseg1	0xFFFF FFFF BFFF FFFF through 0xFFFF FFFF A000 0000	Kernel	Kernel	2^{29} bytes	Always	32-bit Compatibility
	kseg0	0xFFFF FFFF 9FFF FFFF through 0xFFFF FFFF 8000 0000	Kernel	Kernel	2^{29} bytes	Always	32-bit Compatibility
	xkseg	0xFFFF FFFF 7FFF FFFF through 0xC000 0000 0000 0000	Kernel	Kernel	$(2^{SEGBITS} - 2^{31})$ bytes[1]	KX	64-bit
0b10	xkphys	0xBFFF FFFF FFFF FFFF through 0x8000 0000 0000 0000	Kernel	Kernel	8 2^{PABITS} byte[1] regions within the 2^{62} byte Segment	KX	64-bit
0b01	xsseg xksseg	0x7FFF FFFF FFFF FFFF through 0x4000 0000 0000 0000	Supervisor	Supervisor Kernel	$2^{SEGBITS}$ bytes[1]	SX	64-bit
0b00	xuseg xsuseg xkuseg	0x3FFF FFFF FFFF FFFF through 0x0000 0000 8000 0000	User	User Supervisor Kernel	$(2^{SEGBITS} - 2^{31})$ bytes[1]	UX	64-bit
	useg suseg kuseg	0x0000 0000 7FFF FFFF through 0x0000 0000 0000 0000	User	User Supervisor Kernel	2^{31} bytes	Always	32-bit Compatibility

1. See 4.2.2 “Segment and Segment Size (SEGBITS)” and 4.2.3 “Physical Address Size (PABITS)” for an explanation of the symbols *SEGBITS* and *PABITS*, respectively

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.

References to 64-bit Segments (as shown in the “Segment Type” column of Table 4.1) are enabled only if the appropriate 64-bit Address Enable is on (see Section 3.5.1 on page 20, and the “64-bit Enable” column of Table 4.1). References to 32-bit Compatibility Segments are always enabled.

4.4 Compliance

A MIPS64/microMIPS64 compliant processor must implement the following 32-bit Compatibility Segments:

- useg/kuseg
- kseg0
- kseg1

In addition, a MIPS64/microMIPS64 compliant processor using the TLB-based address translation mechanism must also implement the kseg3 32-bit Compatibility 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 64-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.

Table 4.2 Address Space Access and TLB Refill Selection as a Function of Operating Mode

Virtual Address Range		Segment Name(s)	Action when Referenced from Operating Mode		
Symbolic	Assuming *SEGBITS* = 40, *PABITS* = 36		User Mode[1]	Supervisor Mode	Kernel Mode
0xFFFF FFFF FFFF FFFF through 0xFFFF FFFF E000 0000	0xFFFF FFFF FFFF FFFF through 0xFFFF FFFF E000 0000	kseg3	Address Error	Address Error	Mapped Refill Vector: TLB (KX=0) XTLB(KX=1) See Section 4.9 for special behavior when $Debug_{DM}$ = 1
0xFFFF FFFF DFFF FFFF through 0xFFFF FFFF C000 0000	0xFFFF FFFF DFFF FFFF through 0xFFFF FFFF C000 0000	sseg ksseg	Address Error	Mapped Refill Vector[2]: TLB (KX=0) XTLB(KX=1)	Mapped Refill Vector[2]: TLB (KX=0) XTLB(KX=1)

Table 4.2 Address Space Access and TLB Refill Selection as a Function of Operating Mode

Virtual Address Range		Segment Name(s)	Action when Referenced from Operating Mode		
Symbolic	Assuming *SEGBITS* = 40, *PABITS* = 36		User Mode[1]	Supervisor Mode	Kernel Mode
0xFFFF FFFF BFFF FFFF through 0xFFFF FFFF A000 0000	0xFFFF FFFF BFFF FFFF through 0xFFFF FFFF A000 0000	kseg1	Address Error	Address Error	Unmapped, Uncached See Section 4.6
0xFFFF FFFF 9FFF FFFF through 0xFFFF FFFF 8000 0000	0xFFFF FFFF 9FFF FFFF through 0xFFFF FFFF 8000 0000	kseg0	Address Error	Address Error	Unmapped See Section 4.6
0xFFFF FFFF 7FFF FFFF through 0xC000 0000 0000 0000 + $2^{SEGBITS} - 2^{31}$	0xFFFF FFFF 7FFF FFFF through 0xC000 00FF 8000 0000		Address Error	Address Error	Address Error
0xC000 0000 0000 0000 + $2^{SEGBITS} - 2^{31} - 1$ through 0xC000 0000 0000 0000	0xC000 00FF 7FFF FFFF through 0xC000 0000 0000 0000	xkseg	Address Error	Address Error	Address Error if KX = 0 Mapped if KX = 1 Refill Vector: XTLB
0xBFFF FFFF FFFF FFFF through 0x8000 0000 0000 0000	0xBFFF FFFF FFFF FFFF through 0x8000 0000 0000 0000	xkphys	Address Error	Address Error	Address Error if KX = 0 or in certain address ranges within the Segment Unmapped See Section 4.7
0x7FFF FFFF FFFF FFFF through 0x4000 0000 0000 0000 + $2^{SEGBITS}$	0x7FFF FFFF FFFF FFFF through 0x4000 0100 0000 0000		Address Error	Address Error	Address Error
0x4000 0000 0000 0000 + $2^{SEGBITS} - 1$ through 0x4000 0000 0000 0000	0x4000 00FF FFFF FFFF through 0x4000 0000 0000 0000	xsseg xksseg	Address Error	Address Error if SX = 0 Mapped if SX = 1 Refill Vector: XTLB	Address Error if SX = 0 Mapped if SX = 1 Refill Vector: XTLB

Table 4.2 Address Space Access and TLB Refill Selection as a Function of Operating Mode

Virtual Address Range		Segment Name(s)	Action when Referenced from Operating Mode		
Symbolic	Assuming *SEGBITS* = 40, *PABITS* = 36		User Mode[1]	Supervisor Mode	Kernel Mode
0x3FFF FFFF FFFF FFFF through 0x0000 0000 0000 0000 + $2^{SEGBITS}$	0x3FFF FFFF FFFF FFFF through 0x0000 0100 0000 0000		Address Error	Address Error	Address Error
0x0000 0000 0000 0000 + $2^{SEGBITS}$ - 1 through 0x0000 0000 8000 0000	0x0000 00FF FFFF FFFF through 0x0000 0000 8000 0000	xuseg xsuseg xkuseg	Address Error if UX = 0 Mapped if UX = 1 Refill Vector: XTLB	Address Error if UX = 0 Mapped if UX = 1 Refill Vector: XTLB	Address Error if UX = 0 Mapped if UX = 1 Refill Vector: XTLB See Section 4.8 for implementation dependent behavior when $Status_{ERL}$=1
0x0000 0000 7FFF FFFF through 0x0000 0000 0000 0000	0x0000 0000 7FFF FFFF through 0x0000 0000 0000 0000	useg suseg kuseg	Mapped Refill Vector: TLB (UX=0) XTLB(UX=1)	Mapped Refill Vector: TLB (UX=0) XTLB(UX=1)	Unmapped if $Status_{ERL}$=1 See Section 4.8 Mapped if $Status_{ERL}$=0 Refill Vector: TLB (UX=0) XTLB(UX=1)

1. See Section 4.10 for the special treatment of the address for data references when the processor is running in User Mode and the UX bit is zero.
2. Note that the Refill Vector for references to sseg/ksseg is determined by the state of the KX bit, not the SX bit.

4.6 Address Translation and Cacheability & Coherency Attributes for the kseg0 and kseg1 Segments

The kseg0 and kseg1 Unmapped Segments provide a window into the least significant 2^{29} bytes of physical memory, and, as such, are not translated using the TLB or other address translation unit. The cacheability and coherency attribute of the kseg0 Segment is supplied by the K0 field of the CP0 *Config* register. The cacheability and coherency

attribute for the kseg1 Segment is always Uncached. Table 4.3 describes how this transformation is done, and the source of the cacheability and coherency attributes for each Segment.

Table 4.3 Address Translation and Cacheability and Coherency Attributes for the kseg0 and kseg1 Segments

Segment Name	Virtual Address Range	Generates Physical Address	Cache Attribute
kseg1	0xFFFF FFFF BFFF FFFF through 0xFFFF FFFF A000 0000	0x0000 0000 1FFF FFFF through 0x0000 0000 0000 0000	Uncached
kseg0	0xFFFF FFFF 9FFF FFFF through 0xFFFF FFFF 8000 0000	0x0000 0000 1FFF FFFF through 0x0000 0000 0000 0000	From K0 field of *Config* Register

4.7 Address Translation and Cacheability and Coherency Attributes for the xkphys Segment

The xkphys Unmapped Segment is actually composed of 8 address ranges, each of which provides a window into the entire 2^{PABITS} bytes of physical memory and, as such, is not translated using the TLB or other address translation unit. For this Segment, the cacheability and coherency attribute is taken from $VA_{61..59}$ and has the same encoding as that shown in Table 9.9. An Address Error Exception occurs if $VA_{58..PABITS}$ are non-zero. If no Address Error Exception occurs, the physical address is taken from the $VA_{PABITS-1..0}$ virtual address field. Table 4.4 shows the interpretation of the various fields of the virtual address when referencing the xkphys Segment.

Figure 4-2 Address Interpretation for the xkphys Segment

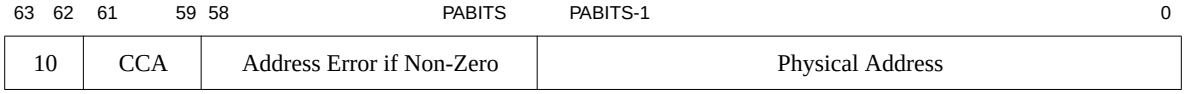

Table 4.4 Address Translation and Cacheability Attributes for the xkphys Segment

Virtual Address Range		Generates Physical Address	Cache Attribute
Symbolic	Assuming *PABITS* = 36		
0xBFFF FFFF FFFF FFFF through 0xB800 0000 0000 0000 + 2^{PABITS}	0xBFFF FFFF FFFF FFFF through 0xB800 0010 0000 0000	Address Error	N/A

Table 4.4 Address Translation and Cacheability Attributes for the xkphys Segment

Virtual Address Range		Generates Physical Address	Cache Attribute
Symbolic	**Assuming *PABITS* = 36**		
0xB800 0000 0000 0000 + 2^{PABITS} - 1 through 0xB800 0000 0000 0000	0xB800 000F FFFF FFFF through 0xB800 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Uses encoding 7 of Table 9.9
0xB7FF FFFF FFFF FFFF through 0xB000 0000 0000 0000 + 2^{PABITS}	0xB7FF FFFF FFFF FFFF through 0xB000 0010 0000 0000	Address Error	N/A
0xB000 0000 0000 0000 + 2^{PABITS} - 1 through 0xB000 0000 0000 0000	0xB000 000F FFFF FFFF through 0xB000 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Uses encoding 6 of Table 9.9
0xAFFF FFFF FFFF FFFF through 0xA800 0000 0000 0000 + 2^{PABITS}	0xAFFF FFFF FFFF FFFF through 0xA800 0010 0000 0000	Address Error	N/A
0xA800 0000 0000 0000 + 2^{PABITS} - 1 through 0xA800 0000 0000 0000	0xA800 000F FFFF FFFF through 0xA800 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Uses encoding 5 of Table 9.9
0xA7FF FFFF FFFF FFFF through 0xA000 0000 0000 0000 + 2^{PABITS}	0xA7FF FFFF FFFF FFFF through 0xA000 0010 0000 0000	Address Error	N/A
0xA000 0000 0000 0000 + 2^{PABITS} - 1 through 0xA000 0000 0000 0000	0xA000 000F FFFF FFFF through 0xA000 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Uses encoding 4 of Table 9.9

Table 4.4 Address Translation and Cacheability Attributes for the xkphys Segment

Virtual Address Range		Generates Physical Address	Cache Attribute
Symbolic	**Assuming *PABITS* = 36**		
0x9FFF FFFF FFFF FFFF through 0x9800 0000 0000 0000 + 2^{PABITS}	0x9FFF FFFF FFFF FFFF through 0x9800 0010 0000 0000	Address Error	N/A
0x9800 0000 0000 0000 + 2^{PABITS} - 1 through 0x9800 0000 0000 0000	0x9800 000F FFFF FFFF through 0x9800 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Cacheable (see encoding 3 of Table 9.9)
0x97FF FFFF FFFF FFFF through 0x9000 0000 0000 0000 + 2^{PABITS}	0x97FF FFFF FFFF FFFF through 0x9000 0010 0000 0000	Address Error	N/A
0x9000 0000 0000 0000 + 2^{PABITS} - 1 through 0x9000 0000 0000 0000	0x9000 000F FFFF FFFF through 0x9000 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Uncached (see encoding 2 of Table 9.9)
0x8FFF FFFF FFFF FFFF through 0x8800 0000 0000 0000 + 2^{PABITS}	0x8FFF FFFF FFFF FFFF through 0x8800 0010 0000 0000	Address Error	N/A
0x8800 0000 0000 0000 + 2^{PABITS} - 1 through 0x8800 0000 0000 0000	0x8800 000F FFFF FFFF through 0x8800 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Uses encoding 1 of Table 9.9
0x87FF FFFF FFFF FFFF through 0x8000 0000 0000 0000 + 2^{PABITS}	0x87FF FFFF FFFF FFFF through 0x8000 0010 0000 0000	Address Error	N/A

Table 4.4 Address Translation and Cacheability Attributes for the xkphys Segment

Virtual Address Range		Generates Physical Address	Cache Attribute
Symbolic	**Assuming *PABITS* = 36**		
0x8000 0000 0000 0000 + 2^{PABITS} - 1 through 0x8000 0000 0000 0000	0x8000 000F FFFF FFFF through 0x8000 0000 0000 0000	0x0000 0000 0000 0000 + 2^{PABITS} - 1 through 0x0000 0000 0000 0000	Uses encoding 0 of Table 9.9

4.8 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.

4.9 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 0xFFFF FFFF FF20 0000 through 0xFFFF FFFF FF3F FFFF, inclusive, as a special memory-mapped region in Debug Mode. A MIPS64/microMIPS64 compliant implementation that also implements EJTAG must:

- explicitly range check the address range as given and not assume that the entire region between 0xFFFF FFFF FF20 0000 and 0xFFFF FFFF 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.

4.10 Special Behavior for Data References in User Mode with $Status_{UX}$ = 0

When the processor is running in User Mode, legal addresses have VA_{31} equal zero, and the 32-bit virtual address is sign-extended (really zero-extended because VA_{31} is zero) into a full 64-bit address. As such, one would expect that the normal address bounds checks on the sign-extended 64-bit address would be sufficient. Unfortunately, there are cases in which a program running on a 32-bit processor can generate a data address that is legal in 32 bits, but which is not appropriately sign-extended into 64-bits. For example, consider the following code example:

```
la  r10, 0x80000000
lw  r10, -4(r10)
```

The results of executing this address calculation on 32-bit and 64-bit processors with UX equal zero is shown below:

32-bit Processor	64-bit Processor
0x8000 0000	0xFFFF FFFF 8000 0000
+0xFFFF FFFC	+0xFFFF FFFF FFFF FFFC

32-bit Processor	64-bit Processor
0x7FFF FFFC	0xFFFF FFFF 7FFF FFFC

On a 32-bit processor, the result of this address calculation results in a valid, useg address. On a 64-bit processor, however, the sign-extended address in the base register is added to the sign-extended displacement as a 64-bit quantity which results in a carry-out of bit 31, producing an address that is not properly sign extended.

To provide backward compatibility with 32-bit User Mode code, MIPS64 compliant processors must implement the following special case for data references (and explicitly *not* for instruction references) when the processor is running in User Mode and the UX bit is zero in the *Status* register:

The effective address calculated by a load, store, or prefetch instruction must be sign extended from bit 31 into bits 63..32 of the full 64-bit address, ignoring the previous contents of bits 63..32 of the address, before the final address is checked for address error exceptions or used to access the TLB or cache. This special-case behavior is not performed for instruction references.

This results in a properly zero-extended address for all legal data addresses (which cleans up the address shown in the example above), and results in a properly sign-extended address for all illegal data addresses (those in which bit 31 is a one). Code running in Debug Mode, Kernel Mode, or Supervisor Mode with the appropriate 64-bit address enable off is prohibited from generating an effective address in which there is a carry-out of bit 31. If such an address is produced, the operation of the instruction generating such an address is **UNPREDICTABLE**.

4.11 TLB-Based Virtual Address Translation[1]

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.

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

4.11.2 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 mapping region specifier (R) and the virtual page number (VPN2 and, in Release 2 and subsequent releases, 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, and in Release 2 and subsequent releases, PFNX), a valid (V) bit, a dirty (D) bit, optionally read-inhibit and execute-inhibit (RI & XI) bits and a cache coherency field (C), whose valid encodings are given in Table 9.9. 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.

1 Refer to A.1 "Fixed Mapping MMU" on page 215 and A.2 "Block Address Translation" on page 219 for descriptions of alternative MMU organizations

In Revision 3 of the architecture, the RI and XI bits were added to the TLB to enable more secure access of memory pages. These bits (along with the Dirty bit) allow the implementation of read-only, write-only, no-execute access policies for mapped pages.

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 and the increase in physical address size from the 36-bit limit in Release 1. Light grey fields denote extensions to the right that are required to support 1KB page sizes. Medium grey fields denote extensions to the left that are required to support larger physical addresses. Neither set of extensions is present in an implementation of Release 1 of the Architecture.

Figure 4.3 Contents of a TLB Entry

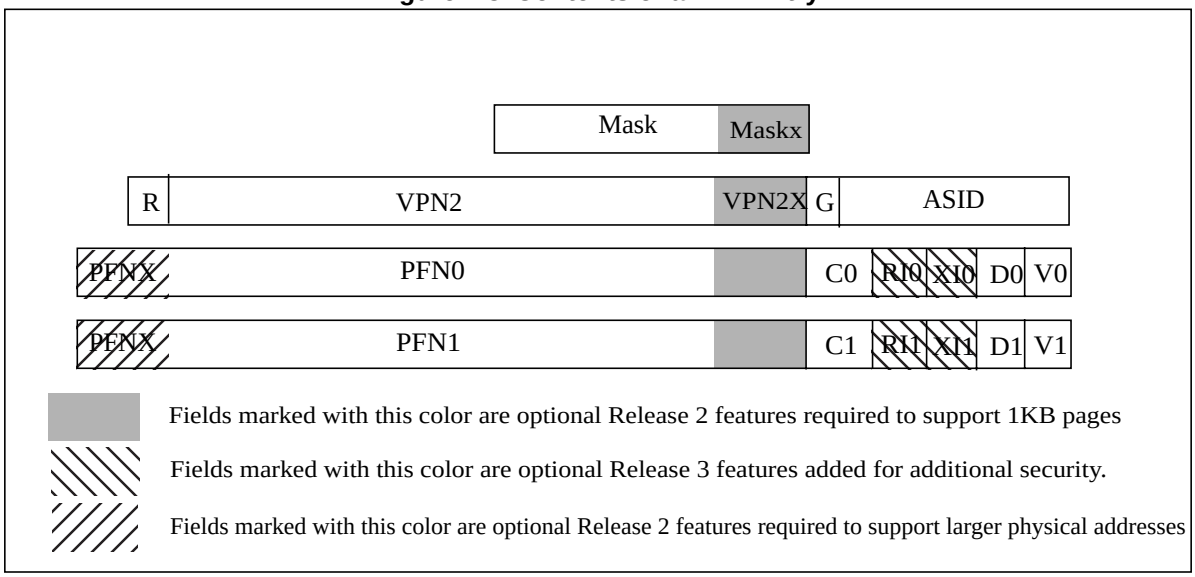

The fields of the TLB entry correspond exactly to the fields in the CP0 *PageMask*, *EntryHi*, *EntryLo0* and *EntryLo1* registers. The even page entries in the TLB (e.g., PFN0) come from *EntryLo0*. Similarly, odd page entries come from *EntryLo1*.

4.11.3 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 (and subsequent releases), processor implentations 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 (and subsequent releases), 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:
 *     a0     =   Maximum TLB index (from MMUSize field of C0_Config1)
 *
 * 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 1 of the
 *        Architecture, and to an ehb in an implementation of Release 2 of
 *        the Architecture. See , "CP0 Hazards," on page 85 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.
 */
   dmtc0  zero, C0_EntryLo0          /* Clear out PFN and valid bits */
   dmtc0  zero, C0_EntryLo1
   mtc0   zero, C0_PageMask          /* Clear out mask register *
/* Start with the base address of kseg0 for the VA part of the TLB */
   la     t0, A_K0BASE               /* A_K0BASE == 0xFFFF.FFFF.8000.0000 */

/*
 * 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:
   dmtc0  t0, C0_EntryHi             /* Write VA candidate */
   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) */
   mfc0   t1, C0_Index               /* Read back flag to check for match */
   bgez   t1, 10b                    /* Branch if about to duplicate an entry */
```

```
   daddiu 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)
 */
   mtc0   a0, C0_Index               /* Use this as next TLB index */
   TLBW_Write_Hazard()               /* Clear Index hazard (ssnop/ehb in R1/2) */
   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
 */
   mtc0   zero, C0_Index
   dmtc0  zero, C0_EntryHi
   jr     ra                         /* Return to caller */
   nop
```

In the code above, 64-bit operations are shown for operations with the TLB. For MIPS64 processors which are running 32-bit software, these instructions may be changed to the corresponding 32-bit instructions.

4.11.4 Address Translation

Release 2 of the Architecture introduced support for 1KB pages, and larger physical addresses. 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 (and subsequent releases) implementations that support 1KB pages concatenate the VPN2 and VPN2X fields to form the virtual page number for a 1KB page
PFN	PFNX || PFN	Release 2 (and subsequent releases) implementations that support larger physical addresses concatenate the PFNX and PFN fields to form the physical page number
Mask	Mask || MaskX	Release 2 (and subsequent releases) 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.
- Bits 63..62 of the virtual address match the region code in the R field of 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 (and optionally RI and XI 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 (and optionally RI and XI 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. If the RI bit is implemented and is set and the reference was a load, a TLB Invalid (or TLBRI) exception is raised. If the XI bit is implemented and is set and the reference was an instruction fetch, a TLB invalid (or TLBXI) exception is raised.

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

1. One used by an implementation of Release 1 of the Architecture, or an implementation of Release 2 (and subsequent releases) of the Architecture which does not include 1KB page support (as denoted by $Config3_{SP}$). This instance is called the “4KB TLB Lookup”.

2. One used by an implementation of Release 2 (and subsequent releases) 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 ← 0
for i in 0...TLBEntries-1
   if (TLB[i]R = va63..62) and
      ((TLB[i]VPN2 and not (TLB[i]Mask)) = (vaSEGBITS-1..13 and not (TLB[i]Mask))) and
      (TLB[i]G or (TLB[i]ASID = EntryHiASID)) 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]Mask
         0b0000 0000 0000 0000: EvenOddBit ← 12 /* 4KB page */
         0b0000 0000 0000 0011: EvenOddBit ← 14 /* 16KB page */
         0b0000 0000 0000 11xx: EvenOddBit ← 16 /* 64KB page */
         0b0000 0000 0011 xxxx: EvenOddBit ← 18 /* 256KB page */
         0b0000 0000 11xx xxxx: EvenOddBit ← 20 /* 1MB page */
         0b0000 0011 xxxx xxxx: EvenOddBit ← 22 /* 4MB page */
         0b0000 11xx xxxx xxxx: EvenOddBit ← 24 /* 16MB page */
         0b0011 xxxx xxxx xxxx: EvenOddBit ← 26 /* 64MB page */
         0b11xx xxxx xxxx xxxx: EvenOddBit ← 28 /* 256MB page */
         otherwise:    UNDEFINED[2]
      endcase
      if vaEvenOddBit = 0 then
         pfn ← TLB[i]PFN0
```

```
            v ← TLB[i]V0
            c ← TLB[i]C0
            d ← TLB[i]D0
            if (Config3RXI or Config3SM) then
                ri ← TLB[i]RI0
                xi ← TLB[i]XI0
            endif
        else
            pfn ← TLB[i]PFN1
            v ← TLB[i]V1
            c ← TLB[i]C1
            d ← TLB[i]D1
            if (Config3RXI or Config3SM) then
                ri ← TLB[i]RI1
                xi ← TLB[i]XI1
            endif
        endif
        if v = 0 then
            SignalException(TLBInvalid, reftype)
        endif
        if (Config3RXI or Config3SM) then
            if (ri = 1) and (reftype = load) then
                if (xi = 0) and (IsPCRelativeLoad(PC))
                    # PC relative loads are allowed where execute is allowed
                else
                    if (PageGrainIEC = 0)
                        SignalException(TLBInvalid, reftype)
                    else
                        SignalException(TLBRI, reftype)
                    endif
                endif
            endif
            if (xi = 1) and (reftype = fetch) then
                if (PageGrainIEC = 0)
                    SignalException(TLBInvalid, reftype)
                else
                    SignalException(TLBXI, reftype)
                endif
            endif
        endif
        if (d = 0) and (reftype = store) then
            SignalException(TLBModified)
        endif
        # pfnPABITS-1-12..0 corresponds to paPABITS-1..12
        pa ← pfnPABITS-1-12..EvenOddBit-12 || vaEvenOddBit-1..0
        found ← 1
        break
    endif
endfor
if found = 0 then
    SignalException(TLBMiss, reftype)
endif
```

2 For brevity, the larger page sizes available through the BigPages feature (1GB and larger) are not shown. The larger page sizes follow the same pattern - 1GB pages would use bit 30 for the EvenOddBit, 4GB would use bit 32. Please refer to Table 4.5 on page 41.

The 1KB TLB Lookup pseudo code is as follows:

```
found ← 0
for i in 0...TLBEntries-1
   if (TLB[i]R = va^63..62) and
      ((TLB[i]VPN2 and not (TLB[i]Mask)) = (vaSEGBITS-1..13 and not (TLB[i]Mask))) and
      (TLB[i]G or (TLB[i]ASID = EntryHiASID)) 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]Mask
         0b0000 0000 0000 0000 00: EvenOddBit ← 10 /* 1KB page */
         0b0000 0000 0000 0000 11: EvenOddBit ← 12 /* 4KB page */
         0b0000 0000 0000 0011 xx: EvenOddBit ← 14 /* 16KB page */
         0b0000 0000 0000 11xx xx: EvenOddBit ← 16 /* 64KB page */
         0b0000 0000 0011 xxxx xx: EvenOddBit ← 18 /* 256KB page */
         0b0000 0000 11xx xxxx xx: EvenOddBit ← 20 /* 1MB page */
         0b0000 0011 xxxx xxxx xx: EvenOddBit ← 22 /* 4MB page */
         0b0000 11xx xxxx xxxx xx: EvenOddBit ← 24 /* 16MB page */
         0b0011 xxxx xxxx xxxx xx: EvenOddBit ← 26 /* 64MB page */
         0b11xx xxxx xxxx xxxx xx: EvenOddBit ← 28 /* 256MB page */
         otherwise:   UNDEFINED[3]
      endcase
      if vaEvenOddBit = 0 then
         pfn ← TLB[i]PFN0
         v ← TLB[i]V0
         c ← TLB[i]C0
         d ← TLB[i]D0
         if (Config3RXI or Config3SM) then
            ri ← TLB[i]RI0
            xi ← TLB[i]XI0
         endif
      else
         pfn ← TLB[i]PFN1
         v ← TLB[i]V1
         c ← TLB[i]C1
         d ← TLB[i]D1
         if (Config3RXI or Config3SM) then
            ri ← TLB[i]RI1
            xi ← TLB[i]XI1
         endif
      endif
      if v = 0 then
         SignalException(TLBInvalid, reftype)
      endif
      if (Config3RXI or Config3SM) then
         if (ri = 1) and (reftype = load) then
            if (xi = 0) and (IsPCRelativeLoad(PC))
               # PC relative loads are allowed where execute is allowed
            else
               if (PageGrainIEC = 0)
```

3 For brevity, the larger page sizes available through the BigPages feature (1GB and larger) are not shown. The larger page sizes follow the same pattern - 1GB pages would use bit 30 for the EvenOddBit, 4GB would use bit 32. Please refer to Table 4.5 on page 41.

```
                    SignalException(TLBInvalid, reftype)
                else
                    SignalException(TLBRI, reftype)
                endif
            endif
        endif
        if (xi = 1) and (reftype = fetch) then
            if (PageGrain_IEC = 0)
                SignalException(TLBInvalid, reftype)
            else
                SignalException(TLBXI, reftype)
            endif
        endif
    endif
    if (d = 0) and (reftype = store) then
        SignalException(TLBModified)
    endif
    # pfn_PABITS-1-10..0 corresponds to pa_PABITS-1..10
    pa ← pfn_PABITS-1-10..EvenOddBit-10 || va_EvenOddBit-1..0
    found ← 1
    break
  endif
endfor
if found = 0 then
  SignalException(TLBMiss, reftype)
endif
```

Table 4.5 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 4.5 indicates which virtual address bit is used to select between the even (EntryLo0) or odd (EntryLo1) entry in the matching TLB entry. The “$PA_{(PABITS-1)..0}$ 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
$PFN_{(PABITS-1)-12..0}$	$PA_{PABITS-1..12}$	Release 1 implementation, or Release 2 (and subsequent releases) implementation without support for 1KB pages
$PFN_{(PABITS-1)-10..0}$	$PA_{PABITS-1..10}$	Release 2 (and subsequent releases) implementation with support for 1KB pages enabled

Table 4.5 Physical Address Generation

Page Size	Even/Odd Select	$PA_{(PABITS-1)..0}$ Generated From:	
		1KB Page Support Unavailable (Release 1) or Disabled (Release 2 & subsequent)	Release 2 (and subsequent) with 1KB Page Support Enabled
1K Bytes	VA_{10}	Not Applicable	$PFN_{(PABITS-1)-10..0} \| VA_{9..0}$

Table 4.5 Physical Address Generation

Page Size	Even/Odd Select	$PA_{(PABITS-1)..0}$ Generated From:	
		1KB Page Support Unavailable (Release 1) or Disabled (Release 2 & subsequent)	Release 2 (and subsequent) with 1KB Page Support Enabled
4K Bytes	VA_{12}	$PFN_{(PABITS-1)-12..0} \| VA_{11..0}$	$PFN_{(PABITS-1)-10..2} \| VA_{11..0}$
16K Bytes	VA_{14}	$PFN_{(PABITS-1)-12..2} \| VA_{13..0}$	$PFN_{(PABITS-1)-10..4} \| VA_{13..0}$
64K Bytes	VA_{16}	$PFN_{(PABITS-1)-12..4} \| VA_{15..0}$	$PFN_{(PABITS-1)-10..6} \| VA_{15..0}$
256K Bytes	VA_{18}	$PFN_{(PABITS-1)-12..6} \| VA_{17..0}$	$PFN_{(PABITS-1)-10..8} \| VA_{17..0}$
1M Bytes	VA_{20}	$PFN_{(PABITS-1)-12..8} \| VA_{19..0}$	$PFN_{(PABITS-1)-10..10} \| VA_{19..0}$
4M Bytes	VA_{22}	$PFN_{(PABITS-1)-12..10} \| VA_{21..0}$	$PFN_{(PABITS-1)-10..12} \| VA_{21..0}$
16M Bytes	VA_{24}	$PFN_{(PABITS-1)-12..12} \| VA_{23..0}$	$PFN_{(PABITS-1)-10..14} \| VA_{23..0}$
64MBytes	VA_{26}	$PFN_{(PABITS-1)-12..14} \| VA_{25..0}$	$PFN_{(PABITS-1)-10..16} \| VA_{25..0}$
256MBytes	VA_{28}	$PFN_{(PABITS-1)-12..16} \| VA_{27..0}$	$PFN_{(PABITS-1)-10..18} \| VA_{27..0}$
1 GBytes[1]	VA_{30}	$PFN_{(PABITS-1)-12..18} \| VA_{29..0}$	$PFN_{(PABITS-1)-10..20} \| VA_{29..0}$
4 GBytes[1]	VA_{32}	$PFN_{(PABITS-1)-12..20} \| VA_{31..0}$	$PFN_{(PABITS-1)-10..22} \| VA_{31..0}$
16 GBytes[1]	VA_{34}	$PFN_{(PABITS-1)-12..22} \| VA_{33..0}$	$PFN_{(PABITS-1)-10..24} \| VA_{33..0}$
64 GByte[1]s	VA_{36}	$PFN_{(PABITS-1)-12..24} \| VA_{35..0}$	$PFN_{(PABITS-1)-10..26} \| VA_{35..0}$
256 GBytes[1]	VA_{38}	$PFN_{(PABITS-1)-12..26} \| VA_{37..0}$	$PFN_{(PABITS-1)-10..28} \| VA_{37..0}$
1 TBytes[1]	VA_{40}	$PFN_{(PABITS-1)-12..28} \| VA_{39..0}$	$PFN_{(PABITS-1)-10..30} \| VA_{39..0}$
4 TBytes[1]	VA_{42}	$PFN_{(PABITS-1)-12..30} \| VA_{41..0}$	$PFN_{(PABITS-1)-10..32} \| VA_{41..0}$
16 TBytes[1]	VA_{44}	$PFN_{(PABITS-1)-12..32} \| VA_{43..0}$	$PFN_{(PABITS-1)-10..34} \| VA_{43..0}$
64 TBytes[1]	VA_{46}	$PFN_{(PABITS-1)-12..34} \| VA_{45..0}$	$PFN_{(PABITS-1)-10..36} \| VA_{45..0}$
256 TBytes[1]	VA_{48}	$PFN_{(PABITS-1)-12..36} \| VA_{47..0}$	$PFN_{(PABITS-1)-10..38} \| VA_{47..0}$

1. .This page size is available only if $Config3_{BPG}$=1.

Chapter 5

Common Device Memory Map

MIPS processors may include memory-mapped IO devices that are packaged as part of the CPU. An example is the Fast Debug Channel, which is a UART-like communication device that uses the EJTAG probe pins to move data to the external world.

The Common Device Memory Map (CDMM) is a region of physical address space that is reserved for mapping IO device configuration registers within a MIPS processor. The CDMM helps aggregate various device mappings into one area, preventing fragmentation of the memory address space. It also enables the use of access control and memory address translation mechanisms for these device registers. The CDMM occupies a maximum of 32KB in the physical address map.

The CMDMM is an optional feature of the architecture. Software detects if CDMM is implemented by reading the $Config3_{CDMM}$ register field (bit 3).

Two blocks are defined for the CDMM -

- *CDMMBase* - A new Coprocessor 0 register that sets the base physical address of the CDMM
- CDMM Access Control and Device Register Block - The 32KB CDMM region is divided into smaller 64-byte aligned blocks called ‘Device Register Blocks’ (DRBs). Each block has access control and status information in access control and status registers (ACSRs), followed by IO device registers.

For implementations that have multiple VPEs, the IO devices and their ACSRs are instantiated once per VPE, but the CDMMBase register is shared among the VPEs.

Implementations are not required to maintain cache coherence for the CDMM region. For that reason, the memory mapped registers located within this region must be accessed only using uncached memory transactions. Accessing these register using a cacheable CCA may result in **UNPREDICTABLE** behavior.

Each of these blocks are now described in detail.

5.1 CDMMBase Register

The physical base address for the CDMM facility is defined by a coprocessor 0 register called *CDMMBase,* (CP0 register 15, select 2). This address must be aligned to a 32KB boundary.

On a 64-bit core with a TLB-based MMU, this region would most likely be mapped to a physical address which can be accessed through one of the kernel unmapped, uncached virtual address segments (kseg1 or xkphys). User-mode access could be allowed through a TLB mapping using an uncached coherency.

On cores that use a FMT MMU, the region would most likely be mapped to the lower 512MB and made accessible via kernel mode. Alternatively, if user-mode access is allowed, this region could be mapped to correspond to the kuseg physical address segment.

On cores that use a BAT MMU, if only kernel mode access is allowed, the region would be mapped to a physical address region reachable through kseg1 or kseg2/3 (using uncached coherency). If user mode access is allowed, the useg BAT entry must use an uncached coherency.

Please refer to Section 9.29 on page 164 for the description of the *CDMMBase* register.

5.2 CDMM - Access Control and Device Register Blocks

The CDMM is divided into 64-byte aligned segments named ‘Device Register Blocks’ (DRBs), Each device occupies at least one DRB. If a device needs additional address space, it can occupy multiple contiguous 64-byte blocks, eg. multiple DRBs which are adjacent in the physical address map. For each device, device type identification and access control information is located in the DRB allocated for the device with the lowest physical address.

Access control information is specified via ‘Access Control and Status Registers’ (ACSRs) that are found at the start of the DRB allocated for the device with the lowest physical address. The ACSR for a device holds the size of the IO device, and hence also act as a pointer to the start of the next device and its’ ACSR. ACSRs are only accessible in kernel mode. The ACSR is followed by the data/control registers for the IO device. Figure 5.1 shows the organization of the CDMM.

Reading any of the IO device registers in either usermode or supervisor mode when such accesses are not allowed, results in all zeros being returned. Writing any of the IO device registers in either usermode or supervisor mode when such accesses are not allowed, results in the write being ignored and the register not being modified. Reading any of the ACSR registers while not in kernel mode results in all zeros being returned. Writing any of the ACSR registers while not in kernel mode results in the write being ignored and the ACSR not being modified.

Since the ACSR act as a pointer that can only increment, the devices must be allocated in the memory space in a specific manner. The first device must be located at the address pointed by the CDMMBase register and any subsequent device is allocated in the next available adjacent DRB.

If the CI bit is set in the CDMMBASE register, the first DRB of the CDMM (at offset 0x0 from the CDMMBase) is reserved for implementation specific use.

Figure 5.1 Example Organization of the CDMM

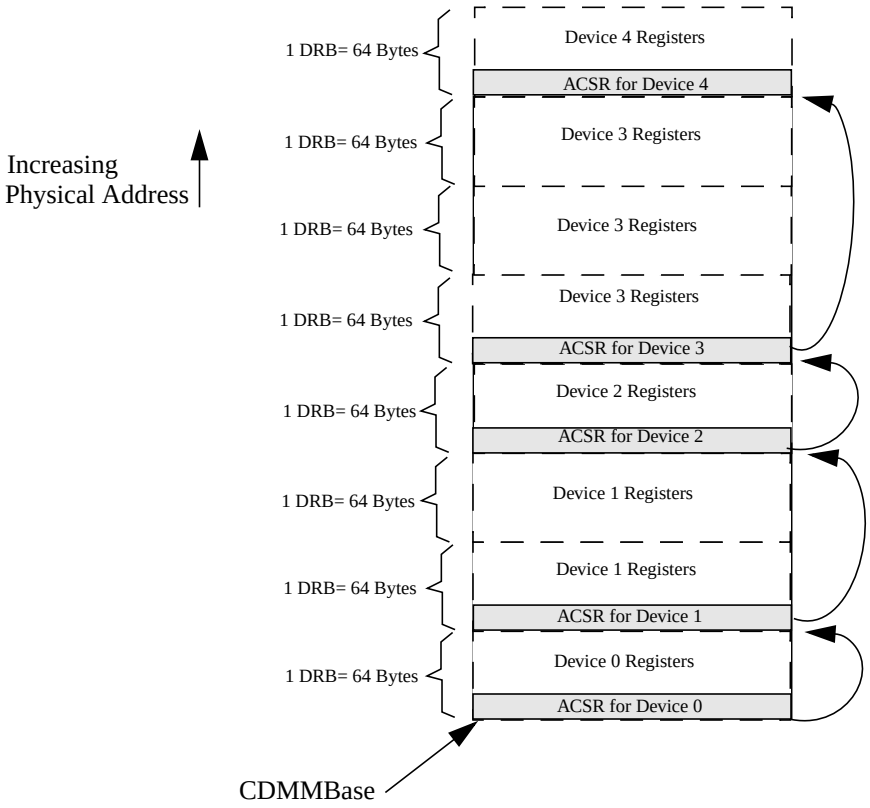

5.2.1 Access Control and Status Registers

The first DRB of a device has 8 bytes of access control address space allocated to it. These 8 bytes can be considered to be two 32-bit registers (on a 32-bit or 64-bit core), or a single 64-bit register (on a 64-bit core). In revision 1.00 of the CDMM, only the lower 32-bits hold access control and status information. The control/status register can be accessed in kernel mode only. Reading this register while not in kernel mode results in all zeros being returned. Writing this register while not in kernel mode results in the write being ignored and the register not being modified.

Figure 5.2 has the format of an Access Control and Status register (shown as a 64-bit register), and Table 5.1 describes the register fields.

Figure 5.2 Access Control and Status Register

63 … 32	31 … 24	23 … 22	21 … 16	15 … 12	11 … 4	3	2	1	0
0	DevType	0	DevSize	DevRev	0	Uw	Ur	Sw	Sr

Table 5.1 Access Control and Status Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
DevType	31:24	This field specifies the type of device. A non-zero value indicates the type of device. A zero value indicates the absence of a device.	R	Preset	Required

Table 5.1 Access Control and Status Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
DevSize	21:16	This field specifies the number of extra 64-byte blocks allocated to this device. A value of 0 indicates that only one 64-byte block is allocated. This also determines the location of the next device block. A device is limited to 4KB of memory.	R	Preset	Required
DevRev	15:12	This field specifies the revision of device. This field is combined with the DevType field to denote the specific device revision.	R	Preset	Required
Uw	3	This bit indicates if user-mode write access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to write to the device while in user mode with access disabled is ignored.	R/W	0	Required
Ur	2	This bit indicates if user-mode read access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to read from the device while in user mode with access disabled is ignored.	R/W	0	Required
Sw	1	This bit indicates if supervisor-mode write access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to write to the device while in supervisor mode with access disabled is ignored.	R/W	0	Required
Sr	0	This bit indicates if supervisor-mode read access to this device is enabled. A value of 1 indicates that access is enabled. A value of 0 indicates that access is disabled. An attempt to read from the device while in supervisor mode with access disabled is ignored.	R/W	0	Required
0	63:32, 11:4	Reserved for future use. Ignored on write; returns zero on read.	R	0	Required

Chapter 6

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 $Status_{BEV}$ 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.
- The addition of an execution hazard sequence which can be used to clear hazards introduced when software writes to a coprocessor 0 register which affects the interrupt system state.

6.1 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 0x180) or the special interrupt vector (0x200), based on the value of $Cause_{IV}$. Software was required to prioritize interrupts as a function of the $Cause_{IP}$ 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.

6.1.1 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 implemented 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 6.1 shows the current interrupt mode of the processor as a function of the coprocessor 0 register fields that can affect the mode.

Table 6.1 Interrupt Modes

$Status_{BEV}$	$Cause_{IV}$	$IntCtl_{VS}$	$Config3_{VINT}$	$Config3_{VEIC}$	Interrupt Mode
1	x	x	x	x	Compatibility
x	0	x	x	x	Compatibility
x	x	=0	x	x	Compatibility
0	1	≠0	1	0	Vectored Interrupt

Table 6.1 Interrupt Modes

$Status_{BEV}$	$Cause_{IV}$	$IntCtl_{VS}$	$Config3_{VINT}$	$Config3_{VEIC}$	Interrupt Mode
0	1	≠0	x	1	External Interrupt Controller
0	1	≠0	0	0	Not Allowed - $IntCtl_{VS}$ is zero if neither Vectored Interrupt nor External Interrupt Controller mode are implemented.
"x" denotes don't care					

6.1.1.1 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 0x180 (if $Cause_{IV}$ = 0) or vector offset 0x200 (if $Cause_{IV}$ = 1). This mode is in effect if any of the following conditions are true:

- $Cause_{IV}$ = 0
- $Status_{BEV}$ = 1
- $IntCtl_{VS}$ = 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 6.2.

Table 6.2 Request for Interrupt Service in Interrupt Compatibility Mode

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

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

```
/*
 * Assumptions:
 *  - Cause_IV = 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 */
   mfc0   k1, C0_Status       /* and Status register for IM bits */
   andi   k0, k0, M_CauseIM   /* Keep only IP bits from Cause */
   and    k0, k0, k1          /* and mask with IM bits */
   beq    k0, zero, Dismiss   /* no bits set - spurious interrupt */
   clz    k0, k0              /* Find first bit set, IP7..IP0; k0 = 16..23 */
   xori   k0, k0, 0x17        /* 16..23 => 7..0 */
   sll    k0, k0, VS          /* Shift to emulate software IntCtl_VS */
   la     k1, VectorBase      /* Get base of 8 interrupt vectors */
   addu   k0, k0, k1          /* Compute target from base and offset */
   jr     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_IM 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.
 */
   eret                       /* Return to interrupted code */

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 */
   dmfc0  k0, C0_EPC          /* Get restart address */
   sd     k0, EPCSave         /* Save in memory */
   mfc0   k0, C0_Status       /* Get Status value */
   sw     k0, StatusSave      /* Save in memory */
   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 */
   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) */
   ld     k1, EPCSave         /*   and EPC */
   mtc0   k0, C0_Status       /* Restore the original value */
   dmtc0  k1, C0_EPC          /*  and EPC */
   /* Restore GPRs and software state */
   eret                       /* Dismiss the interrupt */
```

6.1.1.2 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:

- $Config3_{VInt} = 1$
- $Config3_{VEIC} = 0$
- $IntCtl_{VS} \neq 0$

- $Cause_{IV}$ = 1
- $Status_{BEV}$ = 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 $Cause_{IP}$ bits with the corresponding $Status_{IM}$ bits. If any of these values is 1, and if interrupts are enabled ($Status_{IE}$ = 1, $Status_{EXL}$ = 0, and $Status_{ERL}$ = 0), an interrupt is signaled and a priority encoder scans the values in the order shown in Table 6.3.

Table 6.3 Relative Interrupt Priority for Vectored Interrupt Mode

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

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 6-1.

Figure 6-1 Interrupt Generation for Vectored Interrupt Mode

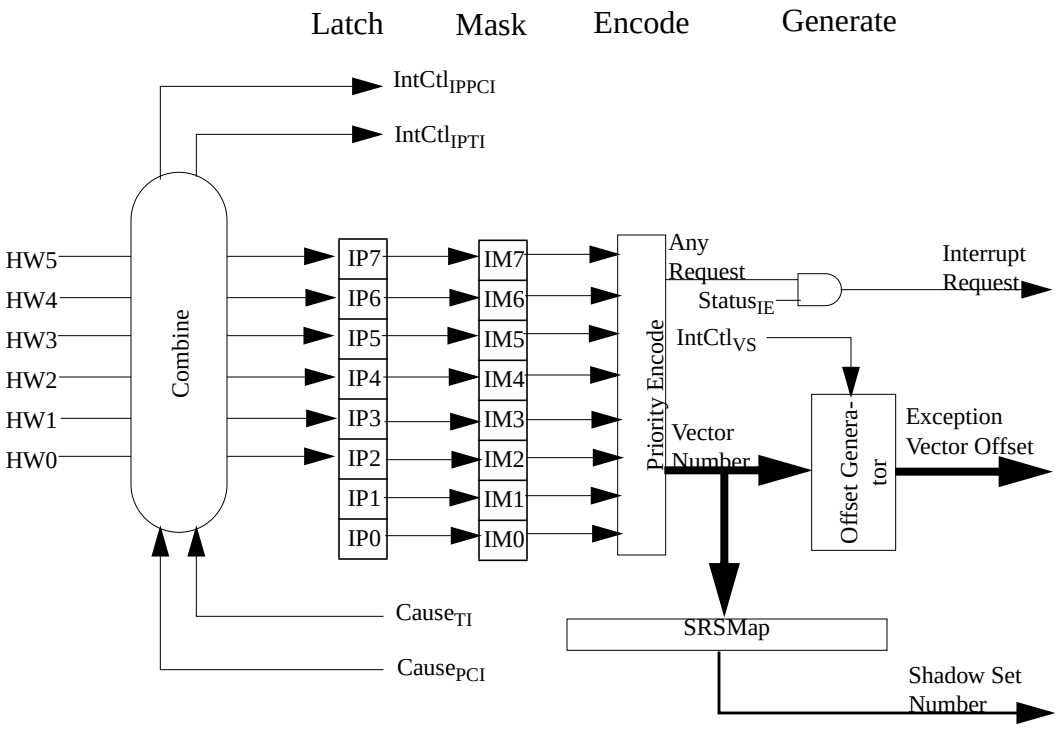

Note that an interrupt request may be deasserted between the time the processor detects the interrupt request and the time that the software interrupt handler runs. The software interrupt handler must be prepared to handle this condition by simply returning from the interrupt via ERET.

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 */
   dmfc0  k0, C0_EPC          /* Get restart address */
   sd     k0, EPCSave         /* Save in memory */
   mfc0   k0, C0_Status       /* Get Status value */
```

```
    sw     k0, StatusSave      /* Save in memory */
    mfc0   k0, C0_SRSCtl       /* Save SRSCtl if changing shadow sets */
    sw     k0, SRSCtlSave
    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 SRSCtl_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 mask, 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 */

/*
 * 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) */
    ld     k1, EPCSave         /*   and EPC */
    mtc0   k0, C0_Status       /* Restore the original value */
    lw     k0, SRSCtlSave      /* Get saved SRSCtl */
    dmtc0  k1, C0_EPC          /*  and EPC */
    mtc0   k0, C0_SRSCtl       /* Restore shadow sets */
    ehb                        /* Clear hazard */
    eret                       /* Dismiss the interrupt */
```

6.1.1.3 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 (and optionally the priority level) of the highest priority interrupt. EIC interrupt mode is in effect if all of the following conditions are true:

- $Config3_{VEIC} = 1$
- $IntCtl_{VS} \neq 0$
- $Cause_{IV} = 1$
- $Status_{BEV} = 0$

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 control-

ler 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 priority level and the vector number of the highest priority interrupt to be serviced. The priority level, 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 lines, which are treated as an encoded value in EIC interrupt mode. There are several implementation options available for the vector offset:

1. The first option is to treat the RIPL value as the vector number for the processor.

2. The second option is to send a separate vector number along with the RIPL to the processor.

3. A third option is to send an entire vector offset along with the RIPL to the processor.

$Status_{IPL}$ (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. Because $Cause_{RIPL}$ is only loaded by the processor when an interrupt exception is signaled, it is available to software during interrupt processing. The vector number that the EIC passes into the core is combined with the $IntCtl_{VS}$ to determine where the interrupt service routines is located. The vector number is not stored in any software visible register. Some implementations may choose to use the RIPL as the vector number, but this is not a requirement.

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 $Cause_{RIPL}$, it also loads the GPR shadow set number into $SRSCtl_{EICSS}$, which is copied to $SRSCtl_{CSS}$ when the interrupt is serviced.

The operation of EIC interrupt mode is shown pictorially in Figure 6-2.

Figure 6-2 Interrupt Generation for External Interrupt Controller Interrupt Mode

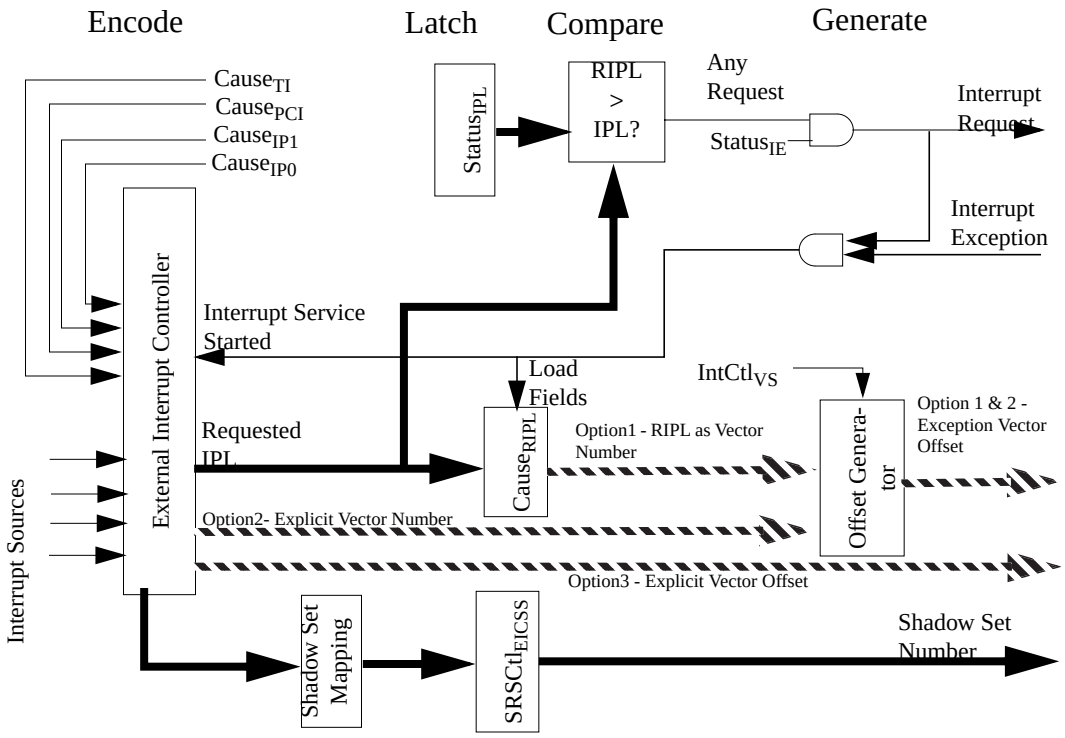

A typical software handler for EIC interrupt mode bypasses the entire sequence of code following the IVexception label shown for the compatibility mode handler above. Instead, the hardware performs the prioritization, dispatching directly to the interrupt processing routine. Unlike the compatibility mode examples, 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$_{RIPL}$ to Status$_{IPL}$ 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 */
   dmfc0  k0, C0_EPC          /* Get restart address */
   srl    k1, k1, S_CauseRIPL /* Right justify RIPL field */
   sd     k0, EPCSave         /* Save in memory */
   mfc0   k0, C0_Status       /* Get Status value */
```

```
    sw     k0, StatusSave      /* Save in memory */
    ins    k0, k1, S_StatusIPL, 6 /* Set IPL to RIPL in copy of Status */
    mfc0   k1, C0_SRSCtl       /* Save SRSCtl if changing shadow sets */
    sw     k1, SRSCtlSave
    /* If switching shadow sets, write new value to SRSCtl_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.
 */
```

6.1.2 Generation of Exception Vector Offsets for Vectored Interrupts

For vectored interrupts (in either VI or EIC interrupt mode - options 1 & 2), 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 0x200 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 6.4 shows the exception vector offset for a representative subset of the vector numbers and values of the $IntCtl_{VS}$ field.

Table 6.4 Exception Vector Offsets for Vectored Interrupts

Vector Number	Value of $IntCtl_{VS}$ Field				
	0b00001	0b00010	0b00100	0b01000	0b10000
0	0x0200	0x0200	0x0200	0x0200	0x0200
1	0x0220	0x0240	0x0280	0x0300	0x0400
2	0x0240	0x0280	0x0300	0x0400	0x0600
3	0x0260	0x02C0	0x0380	0x0500	0x0800
4	0x0280	0x0300	0x0400	0x0600	0x0A00
5	0x02A0	0x0340	0x0480	0x0700	0x0C00
6	0x02C0	0x0380	0x0500	0x0800	0x0E00
7	0x02E0	0x03C0	0x0580	0x0900	0x1000
• • •					

Table 6.4 Exception Vector Offsets for Vectored Interrupts

Vector Number	Value of $IntCtl_{VS}$ Field				
	0b00001	0b00010	0b00100	0b01000	0b10000
61	0x09A0	0x1140	0x2080	0x3F00	0x7C00
62	0x09C0	0x1180	0x2100	0x4000	0x7E00
63	0x09E0	0x11C0	0x2180	0x4100	0x8000

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

$vectorOffset \leftarrow 0x200 + (vectorNumber \times (IntCtl_{VS} \parallel 0b00000))$

6.1.2.1 Software Hazards and the Interrupt System

Software writes to certain coprocessor 0 register fields may change the conditions under which an interrupt is taken. This creates a coprocessor 0 (CP0) hazard, as described in the chapter "CP0 Hazards" on page 85. In Release 1 of the Architecture, there was no architecturally-defined method for bounding the number of instructions which would be executed after the instruction which caused the interrupt state change and before the change to the interrupt state was seen. In Release 2 of the Architecture, the EHB instruction was added, and this instruction can be used by software to clear the hazard.

Table 6.5 lists the CP0 register fields which can cause a change to the interrupt state (either enabling interrupts which were previously disabled or disabling interrupts which were previously enabled).

Table 6.5 Interrupt State Changes Made Visible by EHB

Instruction(s)	CP0 Register Written	CP0 Register Field(s) Modified
MTC0	Status	IM, IPL, ERL, EXL, IE
EI, DI	Status	IE
MTC0	Cause	$IP_{1..0}$
MTC0	PerfCnt Control	IE
MTC0	PerfCnt Counter	Event Count

An EHB, executed after one of these fields is modified by the listed instruction, makes the change to the interrupt state visible no later than the instruction following the EHB.

In the following example, a change to the $Cause_{IM}$ field is made visible by an EHB:

```
mfc0   k0, C0_Status
ins    k0, zero, S_StatusIM4, 1      /* Clear bit 4 of the IM field */
mtc0   k0, C0_Status                 /* Re-write the register */
ehb                                  /* Clear the hazard */
/* Change to the interrupt state is seen no later than this instruction */
```

Similarly, the effects of an DI instruction are made visible by an EHB:

```
di                                   /* Disable interrupts */
```

```
ehb                                  /* Clear the hazard */
/* Change to the interrupt state is seen no later than this instruction */
```

6.2 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.

6.2.1 Exception Priority

Table 6.6 lists all possible exceptions, and the relative priority of each, highest to lowest.

Table 6.6 Priority of Exceptions

Exception	Description	Type
Reset	The Cold Reset signal was asserted to the processor	Asynchronous Reset
Soft Reset	The Reset signal was asserted to the processor	
Debug Single Step	An EJTAG Single Step occurred. Prioritized above other exceptions, including asynchronous exceptions, so that one can single-step into interrupt (or other asynchronous) handlers.	Synchronous Debug
Debug Interrupt	An EJTAG interrupt (EjtagBrk or DINT) was asserted.	Asynchronous Debug
Imprecise Debug Data Break	An imprecise EJTAG data break condition was asserted.	
Nonmaskable Interrupt (NMI)	The NMI signal was asserted to the processor.	Asynchronous
Machine Check	An internal inconsistency was detected by the processor.	
Interrupt	An enabled interrupt occurred.	
Deferred Watch	A watch exception, deferred because EXL was one when the exception was detected, was asserted after EXL went to zero.	
Debug Instruction Break	An EJTAG instruction break condition was asserted. Prioritized above instruction fetch exceptions to allow break on illegal instruction addresses.	Synchronous Debug

Table 6.6 Priority of Exceptions

Exception	Description	Type
Watch - Instruction fetch	A watch address match was detected on an instruction fetch. Prioritized above instruction fetch exceptions to allow watch on illegal instruction addresses.	Synchronous
Address Error - Instruction fetch	A non-word-aligned address was loaded into PC.	
TLB/XTLB Refill - Instruction fetch	A TLB miss occurred on an instruction fetch.	
TLB Invalid - Instruction fetch	The valid bit was zero in the TLB entry mapping the address referenced by an instruction fetch.	
TLB Execute-Inhibit	An instruction fetch matched a valid TLB entry which had the XI bit set.	
Cache Error - Instruction fetch	A cache error occurred on an instruction fetch.	
Bus Error - Instruction fetch	A bus error occurred on an instruction fetch.	
SDBBP	An EJTAG SDBBP instruction was executed.	Synchronous Debug
Instruction Validity Exceptions	An instruction could not be completed because it was not allowed access to the required resources, or was illegal: Coprocessor Unusable, MDMX Unusable, Reserved Instruction. If any two exceptions occur on the same instruction, the Coprocessor Unusable and MDMX Unusable Exceptions take priority over the Reserved Instruction Exception.	Synchronous
Execution Exception	An instruction-based exception occurred: Integer overflow, trap, system call, breakpoint, floating point, coprocessor 2 exception.	
Precise Debug Data Break	A precise EJTAG data break on load/store (address match only) or a data break on store (address+data match) condition was asserted. Prioritized above data fetch exceptions to allow break on illegal data addresses.	Synchronous Debug
Watch - Data access	A watch address match was detected on the address referenced by a load or store. Prioritized above data fetch exceptions to allow watch on illegal data addresses.	Synchronous
Address error - Data access	An unaligned address, or an address that was inaccessible in the current processor mode was referenced, by a load or store instruction	
TLB/XTLB Refill - Data access	A TLB miss occurred on a data access	
TLB Invalid - Data access	The valid bit was zero in the TLB entry mapping the address referenced by a load or store instruction	
TLB Read-Inhibit	A data read access matched a valid TLB entry whose RI bit is set.	
TLB Modified - Data access	The dirty bit was zero in the TLB entry mapping the address referenced by a store instruction	
Cache Error - Data access	A cache error occurred on a load or store data reference	Synchronous or Asynchronous
Bus Error - Data access	A bus error occurred on a load or store data reference	

Table 6.6 Priority of Exceptions

Exception	Description	Type
Precise Debug Data Break	A precise EJTAG data break on load (address+data match only) condition was asserted. Prioritized last because all aspects of the data fetch must complete in order to do data match.	Synchronous Debug

The “Type” column of Table 6.7 describes the type of exception. Table 6.8 explains the characteristics of each exception type.

Table 6.7 Exception Type Characteristics

Exception Type	Characteristics
Asynchronous Reset	Denotes a reset-type exception that occurs asynchronously to instruction execution. These exceptions always have the highest priority to guarantee that the processor can always be placed in a runnable state.
Asynchronous Debug	Denotes an EJTAG debug exception that occurs asynchronously to instruction execution. These exceptions have very high priority with respect to other exceptions because of the desire to enter Debug Mode, even in the presence of other exceptions, both asynchronous and synchronous.
Asynchronous	Denotes any other type of exception that occurs asynchronously to instruction execution. These exceptions are shown with higher priority than synchronous exceptions mainly for notational convenience. If one thinks of asynchronous exceptions as occurring between instructions, they are either the lowest priority relative to the previous instruction, or the highest priority relative to the next instruction. The ordering of the table above considers them in the second way.
Synchronous Debug	Denotes an EJTAG debug exception that occurs as a result of instruction execution, and is reported precisely with respect to the instruction that caused the exception. These exceptions are prioritized above other synchronous exceptions to allow entry to Debug Mode, even in the presence of other exceptions.
Synchronous	Denotes any other exception that occurs as a result of instruction execution, and is reported precisely with respect to the instruction that caused the exception. These exceptions tend to be prioritized below other types of exceptions, but there is a relative priority of synchronous exceptions with each other.

6.2.2 Exception Vector Locations

The Reset, Soft Reset, and NMI exceptions are always vectored to location `0xFFFF.FFFF.BFC0.0000`. EJTAG Debug exceptions are vectored to location `0xFFFF.FFFF.BFC0.0480`, or to location `0xFFFF.FFFF.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 (and subsequent releases), software is allowed to specify the vector base address via the *EBase* register for exceptions that occur when $Status_{BEV}$ equals 0. Table 6.8 gives the vector base address as a function of the exception and whether the *BEV* bit is set in the *Status* register. Table 6.9 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 (and subsequent releases), Table 6.4 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 6.10 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 (and subsequent releases), software must guarantee that $EBase_{15..12}$ 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 0xFFF 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 6.8 Exception Vector Base Addresses

Exception	$Status_{BEV}$	
	0	1
Reset, Soft Reset, NMI	0xFFFF.FFFF.BFC0.0000	
EJTAG Debug (with ProbTrap = 0 in the EJTAG_Control_register)	0xFFFF.FFFF.BFC0.0480	
EJTAG Debug (with ProbTrap = 1 in the EJTAG_Control_register)	0xFFFF.FFFF.FF20.0200	
Cache Error	*For Release 1 of the architecture:* 0xFFFF.FFFF.A000.0000 *For Release 2 of the architecture:* 0xFFFF.FFFF ‖ $EBase_{31..30}$ ‖ 1 ‖ $EBase_{28..12}$ ‖ 0x000 Note that $EBase_{31..30}$ have the fixed value 0b10	0xFFFF.FFFF.BFC0.0200
Other	*For Release 1 of the architecture:* 0xFFFF.FFFF.8000.0000 *For Release 2 of the architecture:* 0xFFFF.FFFF ‖ $EBase_{31..12}$ ‖ 0x000 Note that $EBase_{31..30}$ have the fixed value 0b10	0xFFFF.FFFF.BFC0.0200

Table 6.9 Exception Vector Offsets

Exception	Vector Offset
TLB Refill, EXL = 0	0x000
64-bit XTLB Refill, EXL = 0	0x080
Cache error	0x100
General Exception	0x180

Table 6.9 Exception Vector Offsets

Exception	Vector Offset
Interrupt, $Cause_{IV}$ = 1	0x200 (In Release 2 implementations, this is the base of the vectored interrupt table when $Status_{BEV}$ = 0)
Reset, Soft Reset, NMI	None (Uses Reset Base Address)

Table 6.10 Exception Vectors

Exception	$Status_{BEV}$	$Status_{EXL}$	$Cause_{IV}$	EJTAG ProbTrap	Vector For Release 2 Implementations, assumes that EBase retains its reset state and that $IntCtl_{VS}$ = 0
Reset, Soft Reset, NMI	x	x	x	x	0xFFFF.FFFF.BFC0.0000
EJTAG Debug	x	x	x	0	0xFFFF.FFFF.BFC0.0480
EJTAG Debug	x	x	x	1	0xFFFF.FFFF.FF20.0200
TLB Refill	0	0	x	x	0xFFFF.FFFF.8000.0000
XTLB Refill	0	0	x	x	0xFFFF.FFFF.8000.0080
TLB Refill	0	1	x	x	0xFFFF.FFFF.8000.0180
XTLB Refill	0	1	x	x	0xFFFF.FFFF.8000.0180
TLB Refill	1	0	x	x	0xFFFF.FFFF.BFC0.0200
XTLB Refill	1	0	x	x	0xFFFF.FFFF.BFC0.0280
TLB Refill	1	1	x	x	0xFFFF.FFFF.BFC0.0380
XTLB Refill	1	1	x	x	0xFFFF.FFFF.BFC0.0380
Cache Error	0	x	x	x	0xFFFF.FFFF.A000.0100
Cache Error	1	x	x	x	0xFFFF.FFFF.BFC0.0300
Interrupt	0	0	0	x	0xFFFF.FFFF.8000.0180
Interrupt	0	0	1	x	0xFFFF.FFFF.8000.0200
Interrupt	1	0	0	x	0xFFFF.FFFF.BFC0.0380
Interrupt	1	0	1	x	0xFFFF.FFFF.BFC0.0400
All others	0	x	x	x	0xFFFF.FFFF.8000.0180
All others	1	x	x	x	0xFFFF.FFFF.BFC0.0380
‘x’ denotes don’t care					

6.2.3 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 9.31 on page 151). 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 6.11 shows the value stored in each of the CP0 PC registers, including *EPC*. For implementations of Release 2 of the Architecture if $Status_{BEV}$ = 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.

.

Table 6.11 Value Stored in EPC, ErrorEPC, or DEPC on an Exception

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 63 bits of the address of the instruction, combined with the *ISA Mode* bit
Yes	Yes	Upper 63 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 *ISA Mode* bit

- 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 Status_EXL is 1, all exceptions go through the general exception vector */
/* and neither EPC nor Cause_BD nor SRSCtl are modified */
if Status_EXL = 1 then
   vectorOffset ← 0x180
else
   if InstructionInBranchDelaySlot then
```

```
        EPC ← restartPC/* PC of branch/jump */
        CauseBD ← 1
    else
        EPC ← restartPC              /* PC of instruction */
        CauseBD ← 0
    endif

    /* Compute vector offsets as a function of the type of exception */
    NewShadowSet ← SRSCtlESS         /* Assume exception, Release 2 only */
    if ExceptionType = TLBRefill then
        vectorOffset ← 0x000
    elseif (ExceptionType = XTLBRefill) then
        vectorOffset ← 0x080
    elseif (ExceptionType = Interrupt) then
        if (CauseIV = 0) then
            vectorOffset ← 0x180
        else
            if (StatusBEV = 1) or (IntCtlVS = 0) then
                vectorOffset ← 0x200
            else
                if Config3VEIC = 1 then
                    if (EIC_option1)
                        VecNum ← CauseRIPL
                    elseif (EIC_option2)
                        VecNum ← EIC_VecNum_Signal
                    endif
                    NewShadowSet ← SRSCtlEICSS
                else
                    VecNum ← VIntPriorityEncoder()
                    NewShadowSet ← SRSMapIPL×4+3..IPL×4
                endif
                if (EIC_option3)
                    vectorOffset ← EIC_VectorOffset_Signal
                else
                    vectorOffset ← 0x200 + (VecNum × (IntCtlVS || 0b00000))
                endif
            endif /* if (StatusBEV = 1) or (IntCtlVS = 0) then */
        endif /* if (CauseIV = 0) then */
    endif /* elseif (ExceptionType = Interrupt) then */

    /* Update the shadow set information for an implementation of */
    /* Release 2 of the architecture */
    if (ArchitectureRevision ≥ 2) and (SRSCtlHSS > 0) and (StatusBEV = 0) then
        SRSCtlPSS ← SRSCtlCSS
        SRSCtlCSS ← NewShadowSet
    endif
endif /* if StatusEXL = 1 then */

CauseCE ← FaultingCoprocessorNumber
CauseExcCode ← ExceptionType
StatusEXL ← 1


/* Calculate the vector base address */
if StatusBEV = 1 then
    vectorBase ← 0xFFFF.FFFF.BFC0.0200
else
    if ArchitectureRevision ≥ 2 then
```

```
        /* The fixed value of EBase31..30 forces the base to be in kseg0 or kseg1 */
        vectorBase ← 0xFFFF.FFFF || EBase31..12 || 0x000
    else
        vectorBase ← 0xFFFF.FFFF.8000.0000
    endif
endif

/* Exception PC is the sum of vectorBase and vectorOffset. Vector */
/* offsets > 0xFFF (vectored or EIC interrupts only), require */
/* that EBase15..12 have zeros in each bit position less than or */
/* equal to the most significant bit position of the vector offset */
PC ← vectorBase63..30 || (vectorBase29..0 + vectorOffset29..0)
                          /* No carry between bits 29 and 30 */
```

6.2.4 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

`0xFFFF FFFF BFC0 0480` if the *ProbTrap* bit is zero in the EJTAG_Control_register; `0xFFFF FFFF FF20 0200` if the *ProbTrap* bit is one.

6.2.5 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 6.11. 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 `0xFFFF FFFF BFC0 0000`.

Cause Register ExcCode Value

None

Additional State Saved

None

Entry Vector Used

Reset (0xFFFF FFFF BFC0 0000)

Operation

```
Random ← TLBEntries - 1
EntryLo0PFNX ← 0              # Large physical address implemented
EntryLo1PFNX ← 0              # Large physical address implemented
PageMaskMaskX ← 0             # 1KB page support implemented
PageGrainELPA ← 0             # Large physical address implemented
PageGrainESP ← 0              # 1KB page support implemented
Wired ← 0
HWREna ← 0
EntryHiVPN2X ← 0              # 1KB page support implemented
StatusRP ← 0
StatusBEV ← 1
StatusTS ← 0
StatusSR ← 0
StatusNMI ← 0
StatusERL ← 1
IntCtlVS ← 0
SRSCtlHSS ← HighestImplementedShadowSet
SRSCtlESS ← 0
SRSCtlPSS ← 0
SRSCtlCSS ← 0
SRSMap ← 0
CauseDC ← 0
EBaseExceptionBase ← 0
Config ← ConfigurationState
ConfigK0 ← 2                  # Suggested - see Config register description
Config1 ← ConfigurationState
Config2 ← ConfigurationState
Config3 ← ConfigurationState
WatchLo[n]I ← 0               # For all implemented Watch registers
WatchLo[n]R ← 0               # For all implemented Watch registers
WatchLo[n]W ← 0               # For all implemented Watch registers
PerfCnt.Control[n]IE ← 0      # For all implemented PerfCnt registers
if InstructionInBranchDelaySlot then
   ErrorEPC ← restartPC # PC of branch/jump
else
   ErrorEPC ← restartPC # PC of instruction
endif
PC ← 0xFFFF FFFF BFC0 0000
```

6.2.6 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 copro-

cessor 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 6.11. Note that this value may or may not be predictable.
- PC is loaded with `0xFFFF FFFF BFC0 0000`.

Cause Register ExcCode Value

None

Additional State Saved

None

Entry Vector Used

Reset (`0xFFFF FFFF BFC0 0000`)

Operation

```
EntryLo0PFNX ← 0              # Large physical address implemented
EntryLo1PFNX ← 0              # Large physical address implemented
PageMaskMaskX ← 0             # 1KB page support implemented
PageGrainELPA ← 0             # Large physical address implemented
PageGrainESP ← 0              # 1KB page support implemented
EntryHiVPN2X ← 0              # 1KB page support implemented
ConfigK0 ← 2                  # Suggested - see Config register description
StatusRP ← 0
StatusBEV ← 1
StatusTS ← 0
StatusSR ← 1
StatusNMI ← 0
StatusERL ← 1
WatchLo[n]I ← 0               # For all implemented Watch registers
WatchLo[n]R ← 0               # For all implemented Watch registers
WatchLo[n]W ← 0               # For all implemented Watch registers
PerfCnt.Control[n]IE ← 0      # For all implemented PerfCnt registers
if InstructionInBranchDelaySlot then
   ErrorEPC ← restartPC # PC of branch/jump
else
   ErrorEPC ← restartPC # PC of instruction
endif
PC ← 0xFFFF FFFF BFC0 0000
```

6.2.7 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 6.11.
- PC is loaded with `0xFFFF FFFF BFC0 0000`.

Cause Register ExcCode Value

None

Additional State Saved

None

Entry Vector Used

Reset (`0xFFFF FFFF BFC0 0000`)

Operation

```
Status_BEV ← 1
Status_TS ← 0
Status_SR ← 0
Status_NMI ← 1
Status_ERL ← 1
if InstructionInBranchDelaySlot then
   ErrorEPC ← restartPC # PC of branch/jump
else
   ErrorEPC ← restartPC # PC of instruction
endif
PC ← 0xFFFF FFFF BFC0 0000
```

6.2.8 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 9.32 on page 155)

Additional State Saved

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

Entry Vector Used

General exception vector (offset 0x180)

6.2.9 Address Error Exception

An address error exception occurs under the following circumstances:

- A load or store doubleword instruction is executed in which the address is not aligned on a doubleword boundary.
- 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.
- A reference is made to a a 64-bit address that is outside the range of the 32-bit Compatibility Address Space when 64-bit address references are not enabled.
- A reference is made to an undefined or unimplemented 64-bit address when 64-bit address references are enabled.

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.

Cause Register ExcCode Value

AdEL: Reference was a load or an instruction fetch

AdES: Reference was a store

See Table 9.32 on page 155.

Additional State Saved

Register State	Value
BadVAddr	failing address
$Context_{VPN2}$	UNPREDICTABLE
$XContext_{VPN2}$ $XContext_{R}$	UNPREDICTABLE
$EntryHi_{VPN2}$ $EntryHi_{R}$	UNPREDICTABLE
EntryLo0	UNPREDICTABLE
EntryLo1	UNPREDICTABLE

Entry Vector Used

General exception vector (offset 0x180)

6.2.10 TLB Refill and XTLB Refill Exceptions

A TLB Refill or XTLB 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. Refill exceptions have distinct exception vector offsets: 0x000 for a 32-bit TLB Refill and 0x080 for a 64-bit extended TLB (“XTLB”) refill. The XTLB refill handler is used whenever a reference is made to an enabled 64-bit address space.

Cause Register ExcCode Value

TLBL: Reference was a load or an instruction fetch

TLBS: Reference was a store

See Table 9.32 on page 155.

Additional State Saved

Register State	Value
BadVAddr	Failing address
Context	If $Config3_{CTXTC}$ bit is set, then the bits of the *Context* register corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the bits (starting at bit 31) of the virtual address that missed. If $Config3_{CTXTC}$ bit is clear, then the BadVPN2 field contains $VA_{31..13}$ of the failing address
XContext	If $Config3_{CTXTC}$ bit is set, then the bits of the BadVPN2 field corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the high-order bits (starting at SEGBITS-1) of the virtual address that missed and the R field contains $VA_{63..62}$ of the failing address. If $Config3_{CTXTC}$ bit is clear, then the XContext BadVPN2 field contains $VA_{SEGBITS-1..13}$, and the XContext R field contains $VA_{63..62}$ of the failing address.
EntryHi	The EntryHi VPN2 field contains $VA_{SEGBITS-1..13}$ of the failing address and the EntryHi R field contains $VA_{63..62}$ 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 0x000) if 64-bit addresses are not enabled and $Status_{EXL}$ = 0 at the time of exception.
- XTLB Refill vector (offset 0x080) if 64-bit addresses are enabled and $Status_{EXL}$ = 0 at the time of exception.
- General exception vector (offset 0x180) in either case if $Status_{EXL}$ = 1 at the time of exception

6.2.11 Execute-Inhibit Exception

An Execute-Inhibit exception occurs when the virtual address of an instruction fetch matches a TLB entry whose XI bit is set. This exception type can only occur if the XI bit is implemented within the TLB and is enabled, this is denoted by the $PageGrain_{XIE}$ bit.

Cause Register ExcCode Value

if $PageGrain_{IEC}$ == 0 TLBL

if $PageGrain_{IEC}$ == 1 TLBXI

See Table 9.32 on page 155.

Additional State Saved

Register State	Value
BadVAddr	Failing address
Context	If $Config3_{CTXTC}$ bit is set, then the bits of the *Context* register corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the bits (starting at bit 31) of the virtual address that missed. If $Config3_{CTXTC}$ bit is clear, then the BadVPN2 field contains $VA_{31..13}$ of the failing address
XContext	If $Config3_{CTXTC}$ bit is set, then the bits of the BadVPN2 field corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the high-order bits (starting at SEGBITS-1) of the virtual address that missed and the R field contains $VA_{63..62}$ of the failing address. If $Config3_{CTXTC}$ bit is clear, then the XContext BadVPN2 field contains $VA_{SEGBITS-1..13}$, and the XContext R field contains $VA_{63..62}$ of the failing address.
EntryHi	The EntryHi VPN2 field contains $VA_{SEGBITS-1..13}$ of the failing address and the EntryHi R field contains $VA_{63..62}$ 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 0x180)

6.2.12 Read-Inhibit Exception

An Read-Inhibit exception occurs when the virtual address of a memory load reference matches a TLB entry whose RI bit is set. This exception type can only occur if the RI bit is implemented within the TLB and is enabled, this is denoted by the $PageGrain_{RIE}$ bit. MIPS16 PC-relative loads are a special case and are not affected by the RI bit.

***Cause* Register ExcCode Value**

if $PageGrain_{IEC}$ == 0 TLBL

if $PageGrain_{IEC}$ == 1 TLBRI

See Table 9.32 on page 155.

Additional State Saved

Register State	Value
BadVAddr	Failing address
Context	If $Config3_{CTXTC}$ bit is set, then the bits of the *Context* register corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the bits (starting at bit 31) of the virtual address that missed. If $Config3_{CTXTC}$ bit is clear, then the BadVPN2 field contains $VA_{31..13}$ of the failing address
XContext	If $Config3_{CTXTC}$ bit is set, then the bits of the BadVPN2 field corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the high-order bits (starting at SEGBITS-1) of the virtual address that missed and the R field contains $VA_{63..62}$ of the failing address. If $Config3_{CTXTC}$ bit is clear, then the XContext BadVPN2 field contains $VA_{SEGBITS\text{-}1..13}$, and the XContext R field contains $VA_{63..62}$ of the failing address.
EntryHi	The EntryHi VPN2 field contains $VA_{SEGBITS\text{-}1..13}$ of the failing address, and the EntryHi R field contains $VA_{63..62}$ 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 0x180)

6.2.13 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).

If the RI and XI bits are implemented within the TLB and the $PageGrain_{IEC}$ bit is clear, then this exception also occurs if a valid, matching TLB entry is found with the RI bit set on a memory load reference, or with the XI bit set

on an instruction fetch memory reference. MIPS16 PC-relative loads are a special case and are not affected by the RI bit.

Cause Register ExcCode Value

TLBL: Reference was a load or an instruction fetch

TLBS: Reference was a store

See Table 9.31 on page 151.

Additional State Saved

Register State	Value
BadVAddr	Failing address
Context	If *Config3*$_{CTXTC}$ bit is set, then the bits of the *Context* register corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the bits (starting at bit 31) of the virtual address that missed. If *Config3*$_{CTXTC}$ bit is clear, then the BadVPN2 field contains $VA_{31..13}$ of the failing address
XContext	If *Config3*$_{CTXTC}$ bit is set, then the bits of the BadVPN2 field corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the high-order bits (starting at SEGBITS-1) of the virtual address that missed and the R field contains $VA_{63..62}$ of the failing address. If *Config3*$_{CTXTC}$ bit is clear, then the XContext BadVPN2 field contains $VA_{SEGBITS-1..13}$, and the XContext R field contains $VA_{63..62}$ of the failing address.
EntryHi	The EntryHi VPN2 field contains $VA_{SEGBITS-1..13}$ of the failing address and the EntryHi R field contains $VA_{63..62}$ 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 0x180)

6.2.14 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 9.31 on page 151)

Additional State Saved

Register State	Value
BadVAddr	Failing address
Context	If *Config3*$_{CTXTC}$ bit is set, then the bits of the *Context* register corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the bits (starting at bit 31) of the virtual address that missed. If *Config3*$_{CTXTC}$ bit is clear, then the BadVPN2 field contains $VA_{31..13}$ of the failing address
XContext	If *Config3*$_{CTXTC}$ bit is set, then the bits of the BadVPN2 field corresponding to the set bits of the *VirtualIndex* field of the *ContextConfig* register are loaded with the high-order bits (starting at SEGBITS-1) of the virtual address that missed and the R field contains $VA_{63..62}$ of the failing address. If *Config3*$_{CTXTC}$ bit is clear, then the XContext BadVPN2 field contains $VA_{SEGBITS-1..13}$, and the XContext R field contains $VA_{63..62}$ of the failing address.
EntryHi	The EntryHi VPN2 field contains $VA_{SEGBITS-1..13}$ of the failing address and the EntryHi R field contains $VA_{63..62}$ 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 0x180)

6.2.15 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 0x100)

Operation

```
CacheErr ← ErrorState
Status_ERL ← 1
if InstructionInBranchDelaySlot then
   ErrorEPC ← restartPC # PC of branch/jump
else
   ErrorEPC ← restartPC # PC of instruction
endif
if Status_BEV = 1 then
   PC ← 0xFFFF FFFF BFC0 0200 + 0x100
else
   if ArchitectureRevision ≥ 2 then
      /* The fixed value of EBase_31..30 and bit 29 forced to a 1 puts the */
      /* vector in kseg1 */
      PC ← 0xFFFF.FFFF || EBase_31..30 || 1 || EBase_28..12 || 0x100
   else
      PC ← 0xFFFF FFFF A000 0000 + 0x100
   endif
endif
```

6.2.16 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 9.32 on page 155.

Additional State Saved

None

Entry Vector Used

General exception vector (offset 0x180)

6.2.17 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 9.32 on page 155)

Additional State Saved

None

Entry Vector Used

General exception vector (offset 0x180)

6.2.18 Trap Exception

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

***Cause* Register ExcCode Value**

Tr (See Table 9.32 on page 155)

Additional State Saved

None

Entry Vector Used

General exception vector (offset 0x180)

6.2.19 System Call Exception

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

***Cause* Register ExcCode Value**

Sys (See Table 9.31 on page 151)

Additional State Saved

None

Entry Vector Used

General exception vector (offset 0x180)

6.2.20 Breakpoint Exception

A breakpoint exception occurs when a BREAK instruction is executed.

***Cause* Register ExcCode Value**

Bp (See Table 9.32 on page 155)

Additional State Saved

None

Entry Vector Used

General exception vector (offset 0x180)

6.2.21 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), "⊥" (64-bit) if 64-bit operations are not enabled, or an unimplemented "ε" (ASE).
- An instruction was executed that specifies a *SPECIAL* opcode encoding of the function field that is flagged with "∗" (reserved), "β" (higher-order ISA), or "⊥" (64-bit) if 64-bit operations are not enabled.

- 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 "θ" (partner available), "⊥" (64-bit) if 64-bit operations are not enabled, or an unimplemented "σ" (EJTAG).
- An instruction was executed that specifies a *COPz* opcode encoding of the rs field that is flagged with "∗" (reserved), "β" (higher-order ISA), "⊥" (64-bit) if 64-bit operations are not enabled, 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 when rs is S, D, or W that is flagged with "∗" (reserved), "β" (higher-order ISA), "⊥" (64-bit) if 64-bit operations are not enabled, 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.
- An instruction was executed that specifies a *COP1* opcode encoding when rs is L or PS and 64-bit operations are not enabled, or with a function field encoding 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.
- An instruction was executed that specifies a *COP1X* opcode encoding of the function field that is flagged with "∗" (reserved), or any execution of the *COP1X* opcode when 64-bit operations are not enabled, 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 9.32 on page 155)

Additional State Saved

None

Entry Vector Used

General exception vector (offset 0x180)

6.2.22 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, COP1X, 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. COP2 instructions include MFC2, DMFC2, CFC2, MFHC2, MTC2, DMTC2, CTC2, MTHC2.

Cause Register ExcCode Value

CpU (See Table 9.31 on page 151)

Additional State Saved

Register State	Value
$Cause_{CE}$	unit number of the coprocessor being referenced

Entry Vector Used

General exception vector (offset 0x180)

6.2.23 MDMX Unusable Exception

An MDMX unusable exception occurs if the MDMX instruction is executed and the MX bit of the *Status* register is a 0. Such an exception is used by the operating system to save and restore the state of the MDMX accumulator on a context switch (analogous to the save and restore of the FPRs).

Register ExcCode Value

MDMX (See Table 9.31 on page 151)

Additional State Saved

None

Entry Vector Used

General exception vector (offset 0x180)

6.2.24 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 9.31 on page 151)

Additional State Saved

Register State	Value
FCSR	indicates the cause of the floating point exception

Entry Vector Used

General exception vector (offset 0x180)

6.2.25 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 9.31 on page 151)

Additional State Saved

Defined by the coprocessor

Entry Vector Used

General exception vector (offset 0x180)

6.2.26 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 or SCD instruction does so even if the store would not complete because the *LL* bit is zero.

Register ExcCode Value

WATCH (See Table 9.31 on page 151)

Additional State Saved

Register State	Value
$Cause_{WP}$	indicates that the watch exception was deferred until after both $Status_{EXL}$ and $Status_{ERL}$ 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 0x180)

6.2.27 Interrupt Exception

The interrupt exception occurs when an enabled request for interrupt service is made. See Section 6.1 on page 47 for more information.

Register ExcCode Value

Int (See Table 9.32 on page 155)

Additional State Saved

Register State	Value
$Cause_{IP}$	indicates the interrupts that are pending.

Entry Vector Used

General exception vector (offset 0x180) if the *IV* bit in the *Cause* register is zero.

Interrupt vector (offset 0x200) if the *IV* bit in the *Cause* register is one.

Chapter 7

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, and all shadow sets between 0 and $SRSCtl_{HSS}$, inclusive must be implemented. If this field is zero, only the normal GPRs are implemented.

7.1 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 *SRSCtl* register. When an exception or interrupt occurs, the value of $SRSCtl_{CSS}$ is copied to $SRSCtl_{PSS}$, and $SRSCtl_{CSS}$ is set to the value taken from the appropriate source. On an ERET, the value of $SRSCtl_{PSS}$ is copied back into $SRSCtl_{CSS}$ to restore the shadow set of the mode to which control returns. More precisely, the rules for updating the fields in the *SRSCtl* register on an interrupt or exception are as follows:

1. No field in the *SRSCtl* register is updated if any of the following conditions are true. In this case, steps 2 and 3 are skipped.
   - The exception is one that sets $Status_{ERL}$: NMI or cache error.
   - The exception causes entry into EJTAG Debug Mode
   - $Status_{BEV} = 1$
   - $Status_{EXL} = 1$
2. $SRSCtl_{CSS}$ is copied to $SRSCtl_{PSS}$

3. $SRSCtl_{CSS}$ is updated from one of the following sources:

    - The appropriate field of the *SRSMap* register, based on IPL, if the exception is an interrupt, $Cause_{IV} = 1$, $IntCtl_{VSS} \neq 0$, $Config3_{VEIC} = 0$, and $Config3_{VInt} = 1$. These are the conditions for a vectored interrupt.
    - The EICSS field of the *SRSCtl* register if the exception is an interrupt, $Cause_{IV} = 1$, $IntCtl_{VSS} \neq 0$, 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 $Status_{ERL} = 1$ or $Status_{BEV} = 1$

2. $SRSCtl_{PSS}$ is copied to $SRSCtl_{CSS}$

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 ($Status_{BEV} = 1$).

Privileged software may switch the current shadow set by writing a new value into $SRSCtl_{PSS}$, loading EPC with a target address, and doing an ERET.

7.2 Support Instructions

Table 7.1 Instructions Supporting Shadow Sets

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

Chapter 8

CP0 Hazards

8.1 Introduction

Because resources controlled via Coprocessor 0 affect the operation of various pipeline stages of a MIPS64/microMIPS64 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 MIPS64® 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.

8.2 Types of Hazards

In privileged software, there are two different types of hazards: execution hazards and instruction hazards. Both are defined below.

Implementations using Release 1 of the architecture should refer to their Implementation documentation for the required instruction “spacing” that is required to eliminate these hazards.

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 MIPS64 Release 1 defines the SSNOP instruction to convert instruction issues to cycles in a superscalar design.

8.2.1 Possible Execution Hazards

Execution hazards are those created by the execution of one instruction, and seen by the execution of another instruction. Table 8.1 lists the possible execution hazards that might exist when there are no hardware interlocks.

Table 8.1 Possible Execution Hazards

Producer	→	Consumer	Hazard On
Hazards Related to the TLB			
MTC0	→	TLBR, TLBWI, TLBWR	EntryHi

Table 8.1 Possible Execution Hazards

Producer	→	Consumer	Hazard On
MTC0	→	TLBWI, TLBWR	EntryLo0, EntryLo1, Index, PageMask, PageGrain
MTCO	→	TLBWR	Wired
MTC0	→	TLBP, Load or Store Instruction	$EntryHi_{ASID}$
MTC0	→	Load/store affected by new state	$EntryHi_{ASID}$, WatchHi, WatchLo, Config
TLBP	→	MFC0, TLBWI	Index
TLBR	→	MFC0	EntryHi, EntryLo0, EntryLo1, PageMask
TLBWI, TLBWR	→	TLBP, TLBR, Load/store using new TLB entry	TLB entry
Hazards Related to Exceptions or Interrupts			
MTC0	→	Coprocessor instruction execution depends on the new value of $Status_{CU}$	$Status_{CU}$
MTC0	→	ERET	DEPC, EPC, ErrorEPC, Status
MTC0	→	Interrupted Instruction	$Cause_{IP}$, $Cause_{IV}$ Compare, Count, PerfCnt $Control_{IE}$, PerfCnt Counter, $Status_{IE}$, $Status_{IM}$ EBase SRSCtl SRSMap

Table 8.1 Possible Execution Hazards

Producer	→	Consumer	Hazard On
EI, DI	→	Interrupted Instruction	$Status_{IE}$, $Status_{IM}$
Other Hazards			
LL, LLD	→	MFC0	LLAddr
MTC0	→	CACHE	PageGrain
CACHE	→	MFC0	TagLo
MTC0	→	MFC0	any CoProcessor 0 register

8.2.2 Possible Instruction Hazards

Instruction hazards are those created by the execution of one instruction, and seen by the instruction fetch of another instruction. Table 8.2 lists the possible instruction hazards when there are no hardware interlocks.

Table 8.2 Possible Instruction Hazards

Producer	→	Consumer	Hazard On
Hazards Related to the TLB			
MTC0	→	Instruction fetch seeing the new value	$EntryHi_{ASID}$, WatchHi, WatchLo Config
MTC0	→	Instruction fetch seeing the new value (including a change to ERL followed by an instruction fetch from the useg segment)	Status
TLBWI, TLBWR	→	Instruction fetch using new TLB entry	TLB entry
Hazards Related to Writing the Instruction Stream or Modifying an Instruction Cache Entry			
Instruction stream writes	→	Instruction fetch seeing the new instruction stream	Cache entries
CACHE	→	Instruction fetch seeing the new instruction stream	Cache entries
Other Hazards			
MTC0	→	RDPGPR WRPGPR	$SRSCtl_{PSS}$[1]

1. This is not precisely a hazard on the instruction fetch. Rather it is a hazard on a modification to the previous GPR context field, followed by a previous-context reference to the GPRs. It is considered an instruction hazard rather than an execution hazard because some implementation may require that the previous GPR context be established early in the pipeline, and execution hazards are not meant to cover this case.

8.3 Hazard Clearing Instructions and Events

Table 8.3 lists the instructions designed to eliminate hazards.

Table 8.3 Hazard Clearing Instructions

Mnemonic	Function	Supported Architecture
DERET	Clear both execution and instruction hazards	EJTAG
EHB	Clear execution hazard	Release 2 onwards
ERET	Clear both execution and instruction hazards	All
IRET	Clear both execution and instruction hazards when not chaining to another interrupt.	MCU ASE
JALR.HB	Clear both execution and instruction hazards	Release 2 onwards
JR.HB	Clear both execution and instruction hazards	Release 2 onwards
SSNOP	Superscalar No Operation	Release 1 onwards
SYNCI[1]	Synchronize caches after instruction stream write	Release 2 onwards

1. SYNCI synchronizes caches after an instruction stream write, and before execution of that instruction stream. As such, it is not precisely a coprocessor 0 hazard, but is included here for completeness.

DERET, ERET, and SSNOP are available in Release 1 of the Architecture; EHB, JALR.HB, JR.HB, and SYNCI were added in Release 2 of the Architecture. In both Release 1 and Release 2 of the Architecture, DERET and ERET clear both execution and instruction hazards and they are the only timing-independent instructions which will do this in both releases of the architecture.

Even though DERET and ERET clear hazards between the execution of the instruction and the target instruction stream, an execution hazard may still be created between a write of the *DEPC*, *EPC*, *ErrorEPC*, or *Status* registers and the DERET or ERET instruction.

In addition, an exception or interrupt also clears both execution and instruction hazards between the instruction that created the hazard and the first instruction of the exception or interrupt handler. Said another way, no hazards remain visible by the first instruction of an exception or interrupt handler.

8.3.1 MIPS64 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 MIPS64 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.

8.3.2 microMIPS64 Instruction Encoding

The EHB and SSNOP instructions are encoded using a variant of the NOP encoding. See the EHB and SSNOP instruction description for additional information.

Chapter 9

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.

9.1 Coprocessor 0 Register Summary

Table 9.1 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.

Table 9.1 Coprocessor 0 Registers in Numerical Order

Register Number	Sel[1]	Register Name	Function	Reference	Compliance Level
0	0	Index	Index into the TLB array	Section 9.4 on page 98	Required (TLB MMU); Optional (Others)
0	1	MVPControl	Per-processor register containing global MIPS® MT configuration data	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
0	2	MVPConf0	Per-processor multi-VPE dynamic configuration information	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
0	3	MVPConf1	Per-processor multi-VPE dynamic configuration information	MIPS®MT ASE Specification	Optional
1	0	Random	Randomly generated index into the TLB array	Section 9.5 on page 99	Required (TLB MMU); Optional (Others)
1	1	VPEControl	Per-VPE register containing relatively volatile thread configuration data	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
1	2	VPEConf0	Per-VPE multi-thread configuration information	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
1	3	VPEConf1	Per-VPE multi-thread configuration information	MIPS®MT ASE Specification	Optional
1	4	YQMask	Per-VPE register defining which YIELD qualifier bits may be used without generating an exception	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)

Table 9.1 Coprocessor 0 Registers in Numerical Order

Register Number	Sel[1]	Register Name	Function	Reference	Compliance Level
1	5	VPESchedule	Per-VPE register to manage scheduling of a VPE within a processor	MIPS®MT ASE Specification	Optional
1	6	VPEScheFBack	Per-VPE register to provide scheduling feedback to software	MIPS®MT ASE Specification	Optional
1	7	VPEOpt	Per-VPE register to provide control over optional features, such as cache partitioning control	MIPS®MT ASE Specification	Optional
2	0	EntryLo0	Low-order portion of the TLB entry for even-numbered virtual pages	Section 9.6 on page 100	Required (TLB MMU); Optional (Others)
2	1	TCStatus	Per-TC status information, including copies of thread-specific bits of *Status* and *EntryHi* registers.	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
2	2	TCBind	Per-TC information about TC ID and VPE binding	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
2	3	TCRestart	Per-TC value of restart instruction address for the associated thread of execution	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
2	4	TCHalt	Per-TC register controlling Halt state of TC	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
2	5	TCContext	Per-TC read/write storage for operating system use	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
2	6	TCSchedule	Per-TC register to manage scheduling of a TC	MIPS®MT ASE Specification	Optional
2	7	TCScheFBack	Per-TC register to provide scheduling feedback to software	MIPS®MT ASE Specification	Optional
3	0	EntryLo1	Low-order portion of the TLB entry for odd-numbered virtual pages	Section 9.6 on page 100	Required (TLB MMU); Optional (Others)
3	7	TCOpt	Per-TC register to provide control over optional features, such as cache partitioning control	MIPS®MT ASE Specification	Optional
4	0	Context	Pointer to page table entry in memory	Section 9.7 on page 107	Required (TLB MMU); Optional (Others)
4	1	ContextConfig	Context register configuration	SmartMIPS ASE Specification and Section 9.8 on page 111	Required (SmartMIPS ASE); Optional (Others)

Table 9.1 Coprocessor 0 Registers in Numerical Order

Register Number	Sel[1]	Register Name	Function	Reference	Compliance Level
4	2	UserLocal	User information that can be written by privileged software and read via RDHWR register 29. If the processor implements the MIPS® MT ASE, this is a per-TC register.	Section 9.9 on page 113	Recommended (Release 2)
4	3	XContextConfig	XContext register configuration	Section 9.10 on page 115	Optional
5	0	PageMask	Control for variable page size in TLB entries	Section 9.11 on page 117	Required (TLB MMU); Optional (Others)
5	1	PageGrain	Control for small page support	Section 9.12 on page 121 and SmartMIPS ASE Specification	Required (SmartMIPS ASE); Optional (Release 2)
6	0	Wired	Controls the number of fixed (“wired”) TLB entries	Section 9.13 on page 125	Required (TLB MMU); Optional (Others)
6	1	SRSConf0	Per-VPE register indicating and optionally controlling shadow register set configuration	MIPS®MT ASE Specification	Required (MIPS MT ASE); Optional (Others)
6	2	SRSConf1	Per-VPE register indicating and optionally controlling shadow register set configuration	MIPS®MT ASE Specification	Optional
6	3	SRSConf2	Per-VPE register indicating and optionally controlling shadow register set configuration	MIPS®MT ASE Specification	Optional
6	4	SRSConf3	Per-VPE register indicating and optionally controlling shadow register set configuration	MIPS®MT ASE Specification	Optional
6	5	SRSConf4	Per-VPE register indicating and optionally controlling shadow register set configuration	MIPS®MT ASE Specification	Optional
7	0	HWREna	Enables access via the RDHWR instruction to selected hardware registers	Section 9.14 on page 127	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 9.15 on page 129	Required
9	0	Count	Processor cycle count	Section 9.16 on page 130	Required
9	6-7		Available for implementation dependent user	Section 9.17 on page 130	Implementation Dependent

Table 9.1 Coprocessor 0 Registers in Numerical Order

Register Number	Sel[1]	Register Name	Function	Reference	Compliance Level
10	0	EntryHi	High-order portion of the TLB entry	Section 9.18 on page 131	Required (TLB MMU); Optional (Others)
11	0	Compare	Timer interrupt control	Section 9.19 on page 134	Required
11	6-7		Available for implementation dependent user	Section 9.20 on page 134	Implementation Dependent
12	0	Status	Processor status and control	Section 9.21 on page 135	Required
12	1	IntCtl	Interrupt system status and control	Section 9.22 on page 144	Required (Release 2)
12	2	SRSCtl	Shadow register set status and control	Section 9.23 on page 147	Required (Release 2)
12	3	SRSMap	Shadow set IPL mapping	Section 9.24 on page 150	Required (Release 2 and shadow sets implemented)
12	4	View_IPL	Contiguous view of IM and IPL fields.	MIPS® MCU ASE Specification	Required (MIPS MCU ASE); Optional (Others)
12	5	SRSMap2	Shadow set IPL mapping	MIPS® MCU ASE Specification	Required (MIPS MCU ASE); Optional (Others)
13	0	Cause	Cause of last general exception	Section 9.25 on page 151	Required
13	4	View_RIPL	Contiguous view of IP and RIPL fields.	MIPS® MCU ASE Specification	Required (MIPS MCU ASE); Optional (Others)
14	0	EPC	Program counter at last exception	Section 9.26 on page 157	Required
15	0	PRId	Processor identification and revision	Section 9.27 on page 159	Required
15	1	EBase	Exception vector base register	Section 9.28 on page 161	Required (Release 2)
15	2	CDMMBase	Common Device Memory Map Base register	Section 9.29 on page 164	Optional
15	3	CMGCRBase	Coherency Manager Global Control Register Base register	Section 9.30 on page 166	Optional
16	0	Config	Configuration register	Section 9.31 on page 167	Required

Table 9.1 Coprocessor 0 Registers in Numerical Order

Register Number	Sel[1]	Register Name	Function	Reference	Compliance Level
16	1	Config1	Configuration register 1	Section 9.32 on page 170	Required
16	2	Config2	Configuration register 2	Section 9.33 on page 174	Optional
16	3	Config3	Configuration register 3	Section 9.34 on page 177	Optional
16	3	Config4	Configuration register 4	Section 9.35 on page 183	Optional
16	6-7		Available for implementation dependent user	Section 9.36 on page 187	Implementation Dependent
17	0	LLAddr	Load linked address	Section 9.37 on page 188	Optional
18	0-n	WatchLo	Watchpoint address	Section 9.38 on page 189	Optional
19	0-n	WatchHi	Watchpoint control	Section 9.39 on page 191	Optional
20	0	XContext	Extended Addressing Page Table Context	Section 9.40 on page 193	Required (64-bit TLB MMU) Optional (Others)
21	all		Reserved for future extensions		Reserved
22	all		Available for implementation dependent use	Section 9.41 on page 196	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	UserTraceData1	PDtrace control register	PDtrace Specification	Optional
23	4	TraceIBPC	PDtrace control register	PDtrace Specification	Optional
23	5	TraceDBPC	PDtrace control register	PDtrace Specification	Optional
23	6	Debug2	EJTAG Debug2 register	EJTAG Specification	Optional
24	0	DEPC	Program counter at last EJTAG debug exception	EJTAG Specification	Optional
24	2	TraceContol3	PDtrace control register	PDtrace Specification	Optional

Table 9.1 Coprocessor 0 Registers in Numerical Order

Register Number	Sel[1]	Register Name	Function	Reference	Compliance Level
24	3	UserTraceData2	PDtrace control register	PDtrace Specification	Optional
25	0-n	PerfCnt	Performance counter interface	Section 9.45 on page 200	Recommended
26	0	ErrCtl	Parity/ECC error control and status	Section 9.46 on page 205	Optional
27	0-3	CacheErr	Cache parity error control and status	Section 9.47 on page 206	Optional
28	even selects	TagLo	Low-order portion of cache tag interface	Section 9.48 on page 207	Required (Cache)
28	odd selects	DataLo	Low-order portion of cache data interface	Section 9.49 on page 208	Optional
29	even selects	TagHi	High-order portion of cache tag interface	Section 9.50 on page 209	Required (Cache)
29	odd selects	DataHi	High-order portion of cache data interface	Section 9.51 on page 210	Optional
30	0	ErrorEPC	Program counter at last error	Section 9.52 on page 211	Required
31	0	DESAVE	EJTAG debug exception save register	EJTAG Specification	Optional
31	2-7	KScratch*n*	Scratch Registers for Kernel Mode	Section 9.54 on page 214	Optional; KScratch1 at select 2 and KScratch2 at select 3 are recommended.

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

9.2 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:

Table 9.2 Read/Write Bit Field Notation

Read/Write Notation	Hardware Interpretation	Software Interpretation
R/W	A field in which all bits are readable and writable by software and, potentially, by hardware. Hardware updates of this field are visible by software read. Software updates of this field are visible by hardware read. If the Reset State of this field is “Undefined”, either software or hardware must initialize the value before the first read will return a predictable value. This should not be confused with the formal definition of **UNDEFINED** behavior.	

Table 9.2 Read/Write Bit Field Notation

Read/Write Notation	Hardware Interpretation	Software Interpretation
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 **UNPREDICTABLE** 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 **UNDEFINED** 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.

9.3 Writing CPU Registers

With certain restrictions, software may assume that it can validly write the value read from a coprocessor 0 register back to that register without having unintended side effects. This rule means that software can read a register, modify one field, and write the value back to the register without having to consider the impact of writes to other fields. Processor designers should take this into consideration when using coprocessor 0 register fields that are reserved for implementations and make sure that the use of these bits is consistent with software assumptions.

The most significant exception to this rule is a situation in which the processor modifies the register between the software read and write, such as might occur if an exception or interrupt occurs between the read and write. Software must guarantee that such an event does not occur.

9.4 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 9-1 shows the format of the *Index* register; Table 9.3 describes the *Index* register fields.

Figure 9-1 Index Register Format

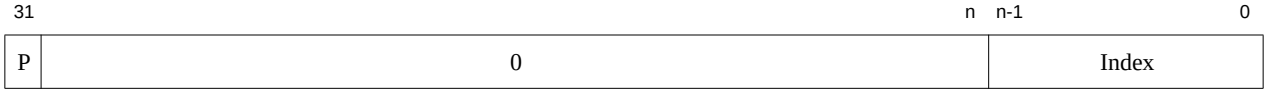

Table 9.3 Index Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
P	31	Probe Failure. Hardware writes this bit during execution of the TLBP instruction to indicate whether a TLB match occurred: Encoding 0: A match occurred, and the *Index* field contains the index of the matching entry Encoding 1: No match occurred and the Index field is **UNPREDICTABLE**	R	Undefined	Required
0	30..n	Must be written as zero; returns zero on read.	0	0	Reserved
Index	n-1..0	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 **UNPREDICTABLE**.	R/W	Undefined	Required

9.5 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 9-2 shows the format of the *Random* register; Table 9.4 describes the *Random* register fields.

Figure 9-2 Random Register Format

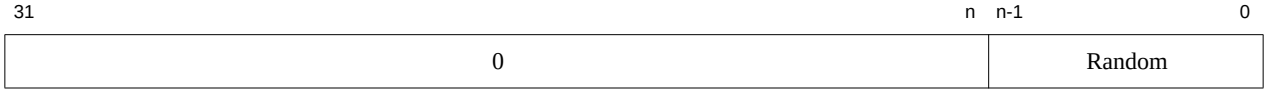

Table 9.4 Random Register Field Descriptions

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

9.6 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/PFNX and Fillfields 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 or DMTC0) do not cause the implicit update of address-related fields in the *BadVAddr* or *Context* registers.

For Release 1 of the Architecture, Figure 9-3 shows the format of the *EntryLo0* and *EntryLo1* registers; Table 9.5 describes the *EntryLo0* and *EntryLo1* register fields.

For Release 2 of the Architecture, Figure 9-4 shows the format of the *EntryLo0* and *EntryLo1* registers; Table 9.6 describes the *EntryLo0* and *EntryLo1* register fields.

For Release 3 of the Architecture, Figure 9-5 shows the format of the *EntryLo0* and *EntryLo1* registers; Figure 9.7 describes the *EntryLo0* and *EntryLo1* register fields.

Figure 9-3 EntryLo0, EntryLo1 Register Format in Release 1 of the Architecture

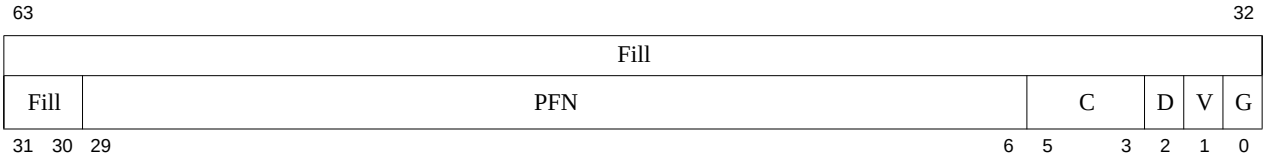

Table 9.5 EntryLo0, EntryLo1 Register Field Descriptions in Release 1 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Fill	63..30	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 9.8 for more information.	R	0	Required
PFN	29..6	Page Frame Number. Corresponds to bits *PABITS*-1..12 of the physical address, where *PABITS* is the width of the physical address in bits. The boundaries of this field change as a function of the value of *PABITS*. See Table 9.8 for more information.	R/W	Undefined	Required

Table 9.5 EntryLo0, EntryLo1 Register Field Descriptions in Release 1 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
C	5..3	Cacheability and Coherency Attribute of the page. See Table 9.9 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
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 9-4 EntryLo0, EntryLo1 Register Format in Release 2 of the Architecture

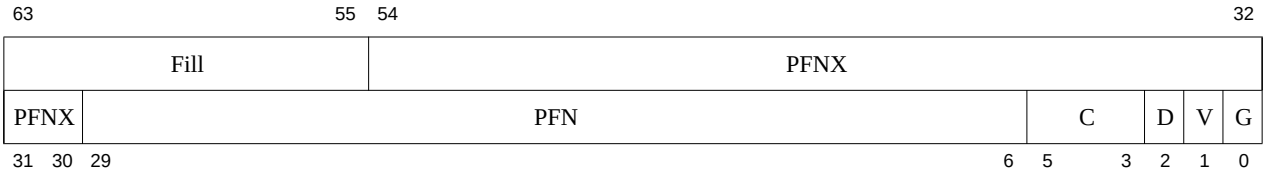

Table 9.6 EntryLo0, EntryLo1 Register Field Descriptions in Release 2 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Fill	63..55	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 9.8 for more information.	R	0	Required

Table 9.6 EntryLo0, EntryLo1 Register Field Descriptions in Release 2 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
PFNX	54..30	Page Frame Number Extension. If the processor is enabled to support large physical addresses ($Config3_{LPA}$ = 1 and $PageGrain_{ELPA}$ = 1), this field is concatenated with the PFN field to form the full page frame number corresponding to the physical address, thereby providing up to 59 bits of physical address. If the processor is enabled to support 1KB pages ($Config3_{SP}$ = 1 and $PageGrain_{ESP}$ = 1), the combined PFNX || PFN fields corresponds to bits *PABITS*-1..10 of the physical address (the field is shifted left by 2 bits relative to the Release 1 definition to make room for $PA_{11..10}$). If the processor is not enabled to support 1KB pages ($Config3_{SP}$ = 0 or $PageGrain_{ESP}$ = 0), the combined PFNX || PFN fields corresponds to 0b00 || bits *PABITS*-1..12 of the physical address (the field is unshifted and the upper two bits must be written as zero). The boundaries of this field change as a function of the value of *PABITS*. See Table 9.8 for more information. If support for large physical addresses is not enabled ($Config3_{LPA}$ = 0 or $PageGrain_{ELPA}$ = 0), these bits are ignored on write and return 0 on read, thereby providing full backward compatibility with implementations of Release 1 of the Architecture.	R/W	0	Optional
PFN	29..6	Page Frame Number. This field contains least-significant bits of the physical page number corresponding to the virtual page. If the processor is enabled to support large physical addresses, the PFNX field, described above is concatenated with the PFN field to form the full page frame number. If the processor is not enabled to support large physical addresses, the entire page frame number is represented by this field. See the description of the PFNX field above for more information. If the processor is enabled to support 1KB pages ($Config3_{SP}$ = 1 and $PageGrain_{ESP}$ = 1), the PFN field corresponds to bits 33..10 of the physical address (the field is shifted left by 2 bits relative to the Release 1 definition to make room for $PA_{11..10}$). If the processor is not enabled to support 1KB pages ($Config3_{SP}$ = 0 or $PageGrain_{ESP}$ = 0), the PFN field corresponds to bits 35..12 of the physical address. The boundaries of this field change as a function of the value of *PABITS*. See Table 9.8 for more information.	R/W	Undefined	Required
C	5..3	The definition of this field is unchanged from Release 1. See Table 9.5 above and Table 9.9 below.	R/W	Undefined	Required
D	2	The definition of this field is unchanged from Release 1. See Table 9.5 above.	R/W	Undefined	Required

Table 9.6 EntryLo0, EntryLo1 Register Field Descriptions in Release 2 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
V	1	The definition of this field is unchanged from Release 1. See Table 9.5 above.	R/W	Undefined	Required
G	0	The definition of this field is unchanged from Release 1. See Table 9.5 above.	R/W	Undefined	Required (TLB MMU)

Figure 9-5 EntryLo0, EntryLo1 Register Format in Release 3 of the Architecture

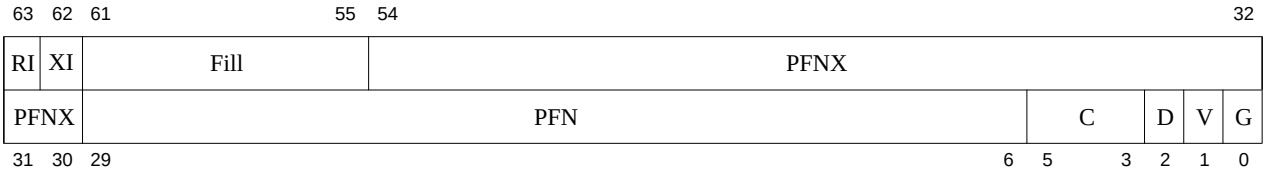

Table 9.7 EntryLo0, EntryLo1 Register Field Descriptions in Release 3 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Fill	61..55	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 9.8 for more information.	R	0	Required if RI and XI fields are not implemented.
RI	63	Read Inhibit. If this bit is set in a TLB entry, an attempt, **other than a MIPS16 PC-relative load,** to read data on the virtual page causes a TLB Invalid or a TLBRI exception, even if the *V* (Valid) bit is set. The *RI* bit is writable only if the *RIE* bit of the *PageGrain* register is set. If the *RIE* bit of *PageGrain* is not set, the *RI* bit of *EntryLo0*/*EntryLo1* is set to zero on any write to the register, regardless of the value written. This bit is optional and its existence is denoted by the $Config3_{RXI}$ or $Config3_{SM}$ register fields.	R/W	0	Required by SmartMIPS ASE; Optional otherwise If not implemented, this bit location is part of the Fill field.
XI	62	Execute Inhibit. If this bit is set in a TLB entry, an attempt to fetch an instruction or to load MIPS16 PC-relative data from the virtual page causes a TLB Invalid or a TLBXI exception, even if the *V* (Valid) bit is set. The *XI* bit is writable only if the *XIE* bit of the *PageGrain* register is set. If the *XIE* bit of *PageGrain* is not set, the *XI* bit of *EntryLo0*/*EntryLo1* is set to zero on any write to the register, regardless of the value written. This bit is optional and its existence is denoted by the $Config3_{RXI}$ or $Config3_{SM}$ register fields.	R/W	0	Required by SmartMIPS ASE; Optional otherwise If not implemented, this bit location is part of the Fill field.

Table 9.7 EntryLo0, EntryLo1 Register Field Descriptions in Release 3 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
PFNX	54..30	Page Frame Number Extension. If the processor is enabled to support large physical addresses ($Config3_{LPA}$ = 1 and $PageGrain_{ELPA}$ = 1), this field is concatenated with the PFN field to form the full page frame number corresponding to the physical address, thereby providing up to 59 bits of physical address. If the processor is enabled to support 1KB pages ($Config3_{SP}$ = 1 and $PageGrain_{ESP}$ = 1), the combined PFNX || PFN fields corresponds to bits *PABITS*-1..10 of the physical address (the field is shifted left by 2 bits relative to the Release 1 definition to make room for $PA_{11..10}$). If the processor is not enabled to support 1KB pages ($Config3_{SP}$ = 0 or $PageGrain_{ESP}$ = 0), the combined PFNX || PFN fields corresponds to 0b00 || bits *PABITS*-1..12 of the physical address (the field is unshifted and the upper two bits must be written as zero). The boundaries of this field change as a function of the value of *PABITS*. See Table 9.8 for more information. If support for large physical addresses is not enabled ($Config3_{LPA}$ = 0 or $PageGrain_{ELPA}$ = 0), these bits are ignored on write and return 0 on read, thereby providing full backward compatibility with implementations of Release 1 of the Architecture.	R/W	0	Optional
PFN	29..6	Page Frame Number. This field contains least-significant bits of the physical page number corresponding to the virtual page. If the processor is enabled to support large physical addresses, the PFNX field, described above is concatenated with the PFN field to form the full page frame number. If the processor is not enabled to support large physical addresses, the entire page frame number is represented by this field. See the description of the PFNX field above for more information. If the processor is enabled to support 1KB pages ($Config3_{SP}$ = 1 and $PageGrain_{ESP}$ = 1), the PFN field corresponds to bits 33..10 of the physical address (the field is shifted left by 2 bits relative to the Release 1 definition to make room for $PA_{11..10}$). If the processor is not enabled to support 1KB pages ($Config3_{SP}$ = 0 or $PageGrain_{ESP}$ = 0), the PFN field corresponds to bits 35..12 of the physical address. The boundaries of this field change as a function of the value of *PABITS*. See Table 9.8 for more information.	R/W	Undefined	Required
C	5..3	The definition of this field is unchanged from Release 1. See Table 9.5 above and Table 9.9 below.	R/W	Undefined	Required
D	2	The definition of this field is unchanged from Release 1. See Table 9.5 above.	R/W	Undefined	Required

Table 9.7 EntryLo0, EntryLo1 Register Field Descriptions in Release 3 of the Architecture

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
V	1	The definition of this field is unchanged from Release 1. See Table 9.5 above.	R/W	Undefined	Required
G	0	The definition of this field is unchanged from Release 1. See Table 9.5 above.	R/W	Undefined	Required (TLB MMU)

Table 9.8 shows the movement of the Fill, PFNX, 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, *PABITS* can never be larger than 36 bits and there is no support for 1KB pages, so only the second row of the table applies to Release 1.

Table 9.8 EntryLo Field Widths as a Function of PABITS

1KB Page Support Enabled?	*PABITS* Value	Corresponding EntryLo Field Bit Ranges			Release 2 Required?
		Fill Field	PFNX Field	PFN Field	
No	59 ≥ *PABITS* > 36	63..(53-(59-*PABITS*)) Example: 63..53 if *PABITS* = 59 63..31 if *PABITS* = 37	(52-(59-*PABITS*))..30 Example: 52..30 if *PABITS* = 59 31..30 if *PABITS* = 37 $EntryLo_{52..30} = PA_{59..36}$	29..6 $EntryLo_{29..6} = PA_{35..12}$	Yes
	36 ≥ *PABITS* > 12	63..(30-(36-*PABITS*)) Example: 63..30 if *PABITS* = 36 63..7 if *PABITS* = 13	Displaced by the Fill Field	(29-(36-*PABITS*))..6 Example: 29..6 if *PABITS* = 36 6..6 if *PABITS* = 13 $EntryLo_{29..6} = PA_{35..12}$	No
Yes	59 ≥ *PABITS* > 34	63..(55-(59-*PABITS*)) Example: 63..55 if *PABITS* = 59 63..31 if *PABITS* = 35	(54-(59-*PABITS*))..30 Example: 54..30 if *PABITS* = 59 31..30 if *PABITS* = 35 $EntryLo_{54..30} = PA_{59..34}$	29..6 $EntryLo_{29..6} = PA_{33..10}$	Yes
	34 ≥ *PABITS* > 10	63..(30-(34-*PABITS*)) Example: 63..30 if *PABITS* = 34 63..7 if *PABITS* = 11	Displaced by the Fill Field	(29-(34-*PABITS*))..6 Example: 29..6 if *PABITS* = 34 6..6 if *PABITS* = 11 $EntryLo_{29..6} = PA_{33..10}$	Yes

Programming Note:

In implementations of Release 2 of the Architecture (and subsequent releases), the PFNX and PFN fields 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 9.9 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 imple-

ment 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 uses a TLB mapping (either for an instruction fetch or for a load/store instruction) which was created with a C field encoding which is RESERVED for the implementation.

Table 9.9 lists the required and optional encodings for the cacheability and coherency attributes.

Table 9.9 Cacheability and Coherency Attributes

C(5:3) Value	Cacheability and 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
5	Available for implementation dependent use	Optional
6	Available for implementation dependent use	Optional
7	Available for implementation dependent use	Optional

9.7 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 is primarily intended for use with the TLB Refill handler, but is also loaded by hardware on an XTLB Refill and may be used by software in that handler. The *Context* register duplicates some of the information provided in the *BadVAddr* register.

If $Config3_{CTXTC}$ =0 and $Config3_{SM}$ =0 then the *Context* register is organized in such a way that the operating system can directly reference a 16-byte structure in memory that describes the mapping. For PTE structures of other sizes, the content of this register can be used by the TLB refill handler after appropriate shifting and masking.

If $Config3_{CTXTC}$ =0 and $Config3_{SM}$ =0 then a TLB exception (TLB Refill, XTLB 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 9-6 shows the format of the *Context* Register when $Config3_{CTXTC}$ =0 and $Config3_{SM}$ =0; Table 9.10 describes the *Context* register fields $Config3_{CTXTC}$ =0 and $Config3_{SM}$ =0.

Figure 9-6 Context Register Format when $Config3_{CTXTC}$=0 and $Config3_{SM}$=0

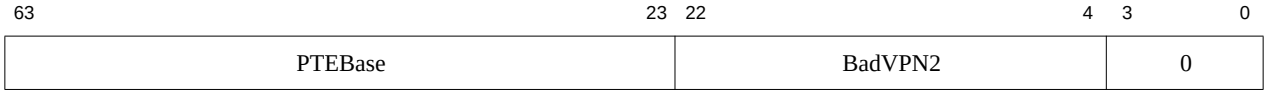

Table 9.10 Context Register Field Descriptions when $Config3_{CTXTC}$=0 and $Config3_{SM}$=0

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
PTEBase	63..23	This field is for use by the operating system and is normally written with a value that allows the operating system to use the *Context* Register as a pointer into the current PTE array in memory.	R/W	Undefined	Required
BadVPN2	22..4	This field is written by hardware on a TLB exception. It contains bits $VA_{31..13}$ of the virtual address that caused the exception.	R	Undefined	Required
0	3..0	Must be written as zero; returns zero on read.	0	0	Reserved

If $Config3_{CTXTC}$ =1 or $Config3_{SM}$ =1 then the pointer implemented by the *Context* register can point to any power-of-two-sized PTE structure within memory. This allows the TLB refill handler to use the pointer without addi-

tional shifting and masking steps. Depending on the value in the *ContextConfig* register, it may point to an 8-byte pair of 32-bit PTEs within a single-level page table scheme, or to a first level page directory entry in a two-level lookup scheme.

If $Config3_{CTXTC}$ =1 or $Config3_{SM}$ =1 then the a TLB exception (Refill, Invalid, or Modified) causes bits $VA_{31:31-((X-Y)-1)}$ to be written to a variable range of bits "(X-1):Y" of the *Context* register, where this range corresponds to the contiguous range of set bits in the *ContextConfig* register. Bits 63:X are R/W to software, and are unaffected by the exception. Bits Y-1:0 are unaffected by the exception. If X = 23 and Y = 4, i.e. bits 22:4 are set in *ContextConfig*, the behavior is identical to the standard MIPS64 *Context* register (bits 22:4 are filled with $VA_{31:13}$). Although the fields have been made variable in size and interpretation, the MIPS64 nomenclature is retained. Bits 63:X are referred to as the *PTEBase* field, and bits X-1:Y are referred to as *BadVPN2*.

If $Config3_{SM}$ =1 then Bits Y-1:0 will always read as 0.

The value of the *Context* register is **UNPREDICTABLE** following a modification of the contents of the *ContextConfig* register.

Figure 9-7 shows the format of the *Context* Register when $Config3_{CTXTC}$ =1 or $Config3_{SM}$ =1; Table 9.11 describes the *Context* register fields $Config3_{CTXTC}$ =1 or $Config3_{SM}$ =1.

Figure 9-7 Context Register Format when $Config3_{CTXTC}$=1 or $Config3_{SM}$=1

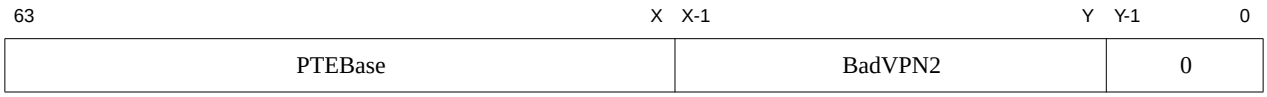

Table 9.11 Context Register Field Descriptions when $Config3_{CTXTC}$=1 or $Config3_{SM}$=1

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
PTEBase	Variable, 63:X where X in {31..0}. May be null.	This field is for use by the operating system and is normally written with a value that allows the operating system to use the *Context* Register as a pointer to an array of data structures in memory corresponding to the address region containing the virtual address which caused the exception.	R/W	Undefined	Required
BadVPN2	Variable, (X-1):Y where X in {32..1} and Y in {31..0}. May be null.	This field is written by hardware on a TLB exception. It contains bits $VA_{31:31-((X-Y)-1)}$ of the virtual address that caused the exception.	R	Undefined	Required
0	Variable, (Y-1):0 where Y in {31:1}. May be null.	Must be written as zero; returns zero on read.	0	0	Reserved

9.8 ContextConfig Register (CP0 Register 4, Select 1)

Compliance Level: *Optional*.

The *ContextConfig* register defines the bits of the *Context* register into which the bits starting from 31 of the virtual address causing a TLB exception will be written, and how many bits of that virtual address will be extracted. Bits above the selected field of the *Context* register are R/W to software and serve as the *PTEBase* field. Bits below the selected field of the *Context* register will be unaffected by TLB exceptions.

The field to contain the virtual address index is defined by a single block of contiguous non-zero bits within the *ContextConfig* register's *VirtualIndex* field. Any zero bits to the right of the least significant one bit cause the corresponding *Context* register bits to be unaffected by TLB exceptions. Any zero bits to the left of the most significant one bit cause the corresponding *Context* register bits to be R/W to software and unaffected by TLB exceptions.

If $Config3_{SM}$ is set, then any zero bits to the right of the least significant one bit causes the corresponding *Context* register bits to be read as zero.

It is permissible to implement a subset of the *ContextConfig* register, in which some number of bits are read-only and set to one or zero as appropriate. It is possible for software to determine which bits are implemented by alternately writing all zeroes and all ones to the register, and reading back the resulting values. All implementations of the *ContextConfig* register must allow for the emulation of the MIPS64/microMIPS64 fixed *Context* register configuration.

This paragraph describes restrictions on how the *ContextConfig register* may be programmed. The set bits of *ContextConfig* define the BadVPN2 field within the *Config* register. The BadVPN2 field cannot contain address bits which are used to index a memory location within the even-odd page pairs used by the JTLB entries. This limits the least significant writeable bit within *ContextConfig* to the bits that represents BadVPN2 of the smallest implemented page size. For example, if the smallest implemented page size is 4KB, virtual address bit 13 is the least significant bit of the BadVPN2 field. This example would restrict the least significant writeable bit within *ContextConfig* to be bit 4 (corresponds to virtual address bit 13) or larger. Another example: if 1KB was the smallest implemented page size then the least significant writeable bit within *ContextConfig* would be bit 2 or larger.

A value of all zeroes means that the full 64 bits of the *Context* register are R/W for software and unaffected by TLB exceptions.

The *ContextConfig* register is optional and its existence is denoted by the $Config3_{CTXTC}$ or $Config3_{SM}$ register fields.

Figure 9.8 shows the formats of the *ContextConfig* Register; Table 9.12 describes the *ContextConfig* register fields.

Figure 9.8 ContextConfig Register Format

Table 9.12 ContextConfig Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
VirtualIndex	31:0	A mask of 0 to 32 contiguous 1 bits in this field causes the corresponding bits of the *Context* register to be written with the bits starting from 31 of the virtual address causing a TLB exception. Behavior of the processor is **UNDEFINED** if non-contiguous 1 bits are written into the register field.	R/W	0x007ffff0	Required

Table 9.13 describes some useful *ContextConfig* values.

Table 9.13 Recommended ContextConfig Values

Value	Page Table Organization	Page Size	PTE Size	Compliance
0x007ffff0	Single Level	4K	64 bits/page	REQUIRED
0x007ffff8	Single Level	2K	32 bits/page	RECOMMENDED

9.9 UserLocal Register (CP0 Register 4, Select 2)

Compliance Level: *Recommended*.

The *UserLocal* register is a read-write register that is not interpreted by the hardware and conditionally readable via the RDHWR instruction.

If the MIPS® MT ASE is implemented, the *UserLocal* register is instantiated per TC.

This register only exists if the $Config3_{ULRI}$ register field is set.

Figure 9-9 shows the format of the *UserLocal* register; Table 9.14 describes the *UserLocal* register fields.

Figure 9-9 UserLocal Register Format

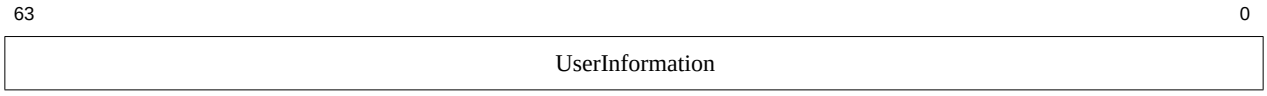

Table 9.14 UserLocal Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
UserInfor-mation	63..0	This field contains software information that is not interpreted by the hardware.	R/W	Undefined	Required

Programming Notes

Privileged software may write this register with arbitrary information and make it accessable to unprivileged software via register 29 (ULR) of the RDHWR instruction. To do so, bit 29 of the *HWREna* register must be set to a 1 to enable unprivileged access to the register. In some operating environments, the *UserLocal* register contains a pointer to a thread-specific storage block that is obtained via the RDHWR register.

9.10 XContextConfig Register (CP0 Register 4, Select 3)

Compliance Level: *Optional*.

The *XContextConfig* register defines the bits of the *XContext* register into which the high order bits (starting at SEGBITS-1) of the virtual address causing a TLB exception will be written, and how many bits of that virtual address will be extracted. Bits above the selected field of the *XContext* register serve as the *PTEBase* and R fields. The *PTEBase* field is R/W to software while the R field is written by hardware. Bits below the selected field of the *Context* register will be unaffected by TLB exceptions.

The field to contain the virtual address index is defined by a single block of contiguous non-zero bits within the *XContextConfig* register's *VirtualIndex* field. Any zero bits to the right of the least significant one bit cause the corresponding *XContext* register bits to be unaffected by hardware. Any zero bits to the left of the most significant one bit designate the location of the R field and cause the remaining *XContext* register bits to be R/W to software and unaffected by TLB exceptions.

It is permissible to implement a subset of the *XContextConfig* register, in which some number of bits are read-only and set to one or zero as appropriate. It is possible for software to determine which bits are implemented by alternately writing all zeroes and all ones to the register, and reading back the resulting values. All implementations of the *XContextConfig* register must allow for the emulation of the MIPS64/microMIPS64 fixed *XContext* register configuration.

This paragraph describes restrictions on how the *XContextConfig register* may be programmed.The set bits of *XContextConfig* define the BadVPN2 field within the *XConfig* register. The BadVPN2 field cannot contain address bits which are used to index a memory location within the even-odd page pairs used by the JTLB entries. This limits the least significant writeable bit within *XContextConfig* to the bits that represents BadVPN2 of the smallest implemented page size. For example, if the smallest implemented page size is 4KB, virtual address bit 13 is the least significant bit of the BadVPN2 field. This example would restrict the least significant writeable bit within *XContextConfig* to be bit 4 (corresponds to virtual address bit 13) or larger. Another example: if 1KB was the smallest implemented page size then the least significant writeable bit within *XContextConfig* would be bit 2 or larger.

In the MIPS64 and microMIPS64 architectures, implementations are allowed to implement virtual address segments which are less than the full 64-bits and have regions in the memory map which are not accessible (accesses to such regions would cause Address Error exceptions). The symbol SEGBITS is used within this document to denote the size of the accessible address segments.The *XConfig* register is meant to be a pointer to a page table data-structure. That page table must reside in memory which is accessible. For that reason, the most significant address bit within the BadVPN2 field can not be larger than the value of SEGBITS-1. This restricts the most significant writeable bit within *XContextConfig* to the bit location that corresponds to $VA_{SEGBITS-1}$ or smaller.

Since the virtual address size is 64-bits, the BadVPN2 field can not extend past bit 63 of the virtual address. This restricts the most significant writeable bit within *XContextConfig* to be bit 54 or smaller.

A value of all zeroes means that the full 64 bits of the *XContext* register are R/W for software and unaffected by TLB exceptions.

The *XContextConfig* register is optional and its existence is denoted by the $Config3_{CTXTC}$ or $Config3_{SM}$ register fields.

The PTEBase fields of *Context* and *XContext* register can be of different width and hold different address pointer values. For this reason, the *XContextConfig* and *ContextConfig* registers must be implemented separately, not sharing any storage.

Figure 9.10 shows the formats of the *XContextConfig* Register; Table 9.15 describes the *XContextConfig* register fields.

Figure 9.10 XContextConfig Register Format

Table 9.15 XContextConfig Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
VirtualIndex	63:0	A mask of 0 to 64 contiguous 1 bits in this field causes the corresponding bits of the *XContext* register to be written with the high-order bits starting at SEGBITS-1 of the virtual address causing a TLB exception. Behavior of the processor is **UNDEFINED** if non-contiguous 1 bits are written into the register field.	R/W	bits SEGBITS-13+3:4 are set (these are the bits corresponding to $VA_{SEGBITS-1:13}$)	Required

9.11 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 9.17. Figure 9-11 shows the format of the *PageMask* register; Table 9.16 describes the *PageMask* register fields.

Figure 9-11 PageMask Register Format if $Config_{BPG}$=0

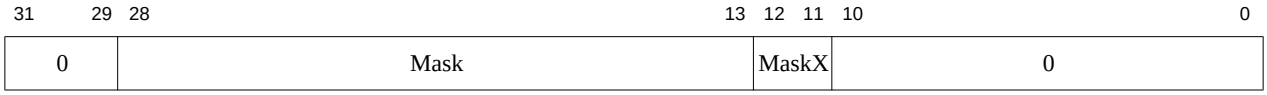

Figure 9-12 PageMask Register Format if $Config_{BPG}$=1

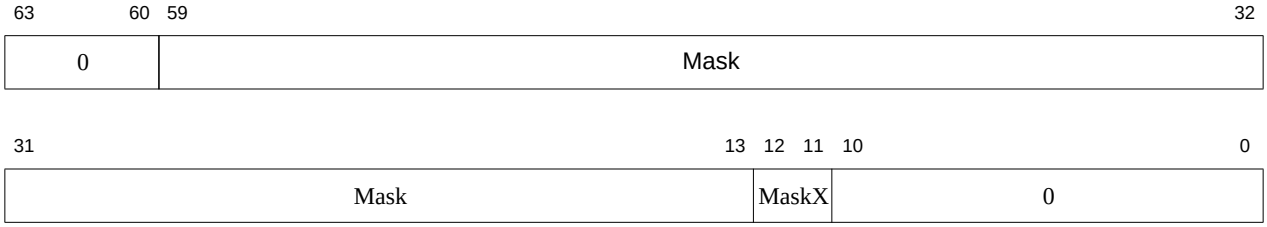

Table 9.16 PageMask Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Mask	if $Config_{BPG}$=0 28..13 if $Config_{BPG}$=1 59..13	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

Table 9.16 PageMask Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
MaskX	12..11	In Release 2 of the Architecture (and subsequent releases), 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 ($Config3_{SP}$ = 1 and $PageGrain_{ESP}$ = 1), these bits are writable and readable, and their values are copied to and from the TLB entry on a TLB write or read, respectively. If 1KB pages are not enabled ($Config3_{SP}$ = 0 or $PageGrain_{ESP}$ = 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 0b11. 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	if $Config_{BPG}$=0 31..29, if $Config_{BPG}$=1 63..60, 10..0	Ignored on write; returns zero on read.	R	0	Required

Table 9.17 Values for the Mask and MaskX[1] Fields of the PageMask Register

Page Size	Values for Mask field (lsb of value is located at $PageMask_{13}$)	Values for MaskX[1] field
1 KByte	0x0	0x0
4 KByte	0x0	0x3
16 KByte	0x3	0x3
64 KByte	0xF	0x3
256 KByte	0x3F	0x3

Table 9.17 Values for the Mask and MaskX[1] Fields of the PageMask Register

Page Size	Values for Mask field (lsb of value is located at $PageMask_{13}$)	Values for MaskX[1] field
1 MByte	0xFF	0x3
4 MByte	0x3FF	0x3
16 MByte	0xFFF	0x3
64 MByte	0x3FFF	0x3
256 MByte	0xFFFF	0x3
1 GByte[2]	0x3FFFF	0x3
4 GByte[2]	0xFFFFF	0x3
16 GByte[2]	0x3FFFFF	0x3
64GByte[2]	0xFFFFFF	0x3
256 GByte[2]	0x3FFFFFF	0x3
1 TByte[2]	0xFFFFFFF	0x3
4 TByte[2]	0x3FFFFFFF	0x3
16 TByte[2]	0xFFFFFFFF	0x3
64 TByte[2]	0x3FFFFFFFF	0x3
256 TByte[2]	0xFFFFFFFFF	0x3

1. $PageMask_{12..11}$ = $PageMask_{MaskX}$ exists only on implementations of Release 2 of the architecture and are treated as if they had the value 0b11 if 1K pages are not enabled ($Config3_{SP}$ = 0 or $PageGrain_{ESP}$ = 0).
2. .This page size is available only if $Config3_{BPG}$=1. The left-most bits of the Mask field necessary to represent this page size exist only if $Config3_{BPG}$=1

It is implementation dependent how many of the encodings described in Table 9.17are 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 9.17, even if the hardware returns a different value on read. Hardware may depend on this requirement in implementing hard-

ware structures

Programming Note:

In implementations of Release 2 (and subsequent releases) of the Architecture, the *MaskX* field of the *PageMask* register must be written with 0b11 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.

9.12 PageGrain Register (CP0 Register 5, Select 1)

Compliance Level: *Required* for implementations of Release 2 (and subsequent releases) of the Architecture that include TLB-based MMUs and support 1KB pages, the XI/RI TLB protection bits, or large physical addresses; *Required* for SmartMIPS™ ASE; otherwise *Optional*.

The *PageGrain* register is a read/write register used for enabling 1KB page support, the XI/RI TLB protection bits, and large physical address support. The *PageGrain* register is present in both the SmartMIPS™ ASE, and in Release 2 (and subsequent releases) of the Architecture. 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™ ASE, the ASE definitions take precedence. Figure 9-13 shows the format of the *PageGrain* register; Table 9.18 describes the *PageGrain* register fields.

Figure 9-13 PageGrain Register Format

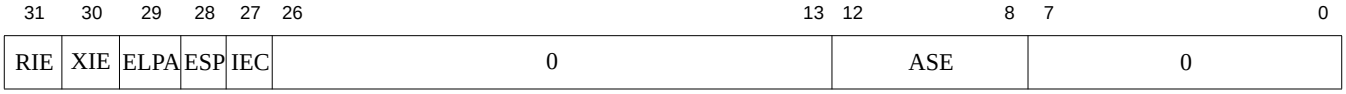

Table 9.18 PageGrain Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
RIE	31	Read Inhibit Enable. Encoding / Meaning: 0 — *RI* bit of the *EntryLo0* and *EntryLo1* registers is disabled and not writeable by software. 1 — *RI* bit of the *EntryLo0* and *EntryLo1* registers is enabled. This bit is optional. The existence of this bit is denoted by either the SM or RXI bits within the *Config3* register. If this bit is not settable then the RI bit within the *EntryLo** registers is not implemented.	R/W or R	0	Required by SmartMIPS ASE; Optional otherwise

Table 9.18 PageGrain Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>XIE</td><td>30</td><td>Execute Inhibit Enable.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>XI bit of the EntryLo0 and EntryLo1 registers is disabled and not writeable by software.</td></tr><tr><td>1</td><td>XI bit of the EntryLo0 and EntryLo1 registers is enabled.</td></tr></table>
This bit is optional. The existence of this bit is denoted by either the SM or RXI bits within the Config3 register. If this bit is not settable then the XI bit within the EntryLo* registers is not implemented.</td><td>R/W
or
R</td><td>0</td><td>Required by SmartMIPS ASE; Optional otherwise</td></tr>
<tr><td>ASE</td><td>12..8</td><td>These fields are control features of the SmartMIPS™ ASE and are not used in implementations of Release 2 of the Architecture unless such an implementation also implements the SmartMIPS™ ASE.</td><td>0</td><td>0</td><td>Required</td></tr>
<tr><td>ELPA</td><td>29</td><td>Enables support for large physical addresses.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Large physical address support is not enabled</td></tr><tr><td>1</td><td>Large physical address support is enabled</td></tr></table>
If this bit is a 1, the following changes occur to coprocessor 0 registers:
• The PFNX field of the EntryLo0 and EntryLo1 registers is writable and concatenated with the PFN field to form the full page frame number.
If $Config3_{LPA}$ = 0, large physical addresses are not implemented, and this bit is ignored on write and returns zero on read.</td><td>R/W</td><td>0</td><td>Required</td></tr>
</table>

Table 9.18 PageGrain Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
ESP	28	Enables support for 1KB pages. Encoding: 0 — 1KB page support is not enabled; 1 — 1KB page support is enabled If this bit is a 1, the following changes occur to coprocessor 0 registers: • The PFN and PFNX fields of the *EntryLo0* and *EntryLo1* registers hold the physical address down to bit 10 (the field is shifted left by 2 bits from the Release 1 definition) • The MaskX field of the *PageMask* register is writable and is concatenated to the right of the Mask field to form the "don't care" mask for the TLB entry. • The VPN2X field of the *EntryHi* register is writable and bits 12..11 of the virtual address. • The virtual address translation algorithm is modified to reflect the smaller page size. If $Config3_{SP}$ = 0, 1KB pages are not implemented, and this bit is ignored on write and returns zero on read.	R/W	0	Required
IEC	27	Enables unique exception codes for the Read-Inhibit and Execute-Inhibit exceptions. Encoding: 0 — Read-Inhbit and Execute-Inhibit exceptions both use the TLBL exception code.; 1 — Read-Inhibit exceptions use the TLBRI exception code. Execute-Inhibit exceptions use the TLBXI exception code For implementations which follow the SmartMIPS ASE, this bit is ignored by the hardware, meaning the Read-Inhibit and Execute-Inhibit exceptions can only use the TLBL exception code.	R/W	0	Required
0	26..13, 7..0	Must be written as zero; returns zero on read.	0	0	Reserved

Programming Note:

In implementations of Release 2 (and subsequent releases) 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 **UNDEFINED** if this sequence is not done.

Field	Required Value
$EntryLo0_{PFN}$, $EntryLo1_{PFN}$	0

Field	Required Value
$EntryLo0_{PFNX}$, $EntryLo1_{PFNX}$	0
$PageMask_{MaskX}$	0b11
$EntryHi_{VPN2X}$	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.

9.13 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 9-14.

Figure 9-14 Wired And Random Entries In The TLB

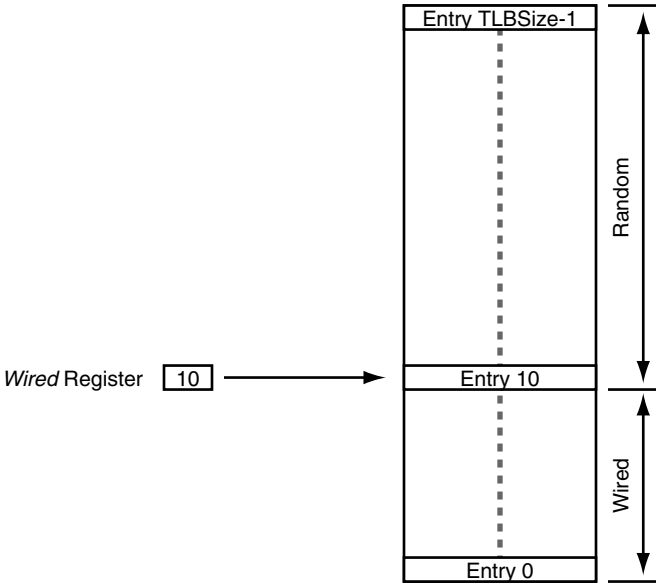

The width of the *Wired* field is calculated in the same manner as that described for the *Index* register. *Wired* entries are fixed, non-replaceable entries which are not overwritten by a TLBWR instruction.*Wired* entries can be overwritten by a TLBWI instruction.

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 9-14 shows the format of the *Wired* register; Table 9.19 describes the *Wired* register fields.

Figure 9-15 Wired Register Format

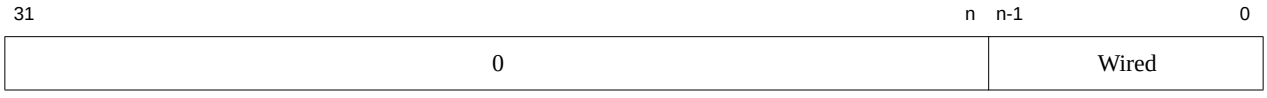

Table 9.19 Wired Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
0	31..n	Must be written as zero; returns zero on read.	0	0	Reserved
Wired	n-1..0	TLB wired boundary	R/W	0	Required

9.14 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 when that instruction is executed in a mode in which coprocessor 0 is not enabled.

Figure 9-16 shows the format of the *HWREna* Register; Table 9.20 describes the *HWREna* register fields.

Figure 9-16 HWREna Register Format

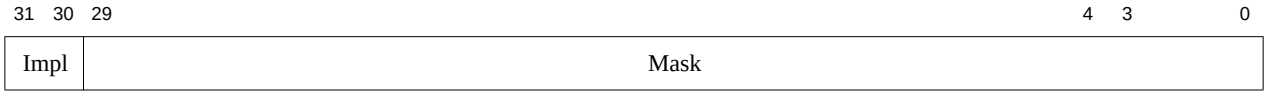

Table 9.20 HWREna Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
31..30	Impl	These bits enable access to the implementation-dependent hardware registers 31 and 30. If a register is not implemented, the corresponding bit returns a zero and is ignored on write. If a register is implemented, access to that register is enabled if the corresponding bit in this field is a 1 and disabled if the corresponding bit is a 0.	R/W	0	Optional - Reserved for Implementations
Mask	29..0	Each bit in this field enables access by the RDHWR instruction to a particular hardware register (which may not be an actual register). If RDHWR register ‘n’ is not implemented, bit ‘n’ of this field returns a zero and is ignored on a write. If RDHWR register ‘n’ is implemented, access to the register is enabled if bit ‘n’ in this field is a 1 and disabled if bit ‘n’ of this field is a 0. See the RDHWR instruction for a list of valid hardware registers. Table 9.21 lists the RDHWR registers, and register number ‘n’ corresponds to bit ‘n’ in this field.	R/W	0	Required

Table 9.21 RDHWR Register Numbers

<table>
<tr><th>Register Number</th><th>Mnemonic</th><th>Description</th><th>Compliance</th></tr>
<tr><td>0</td><td>CPUNum</td><td>Number of the CPU on which the program is currently running. This register provides read access to the coprocessor 0 EBase_{CPUNum} field.</td><td>Required</td></tr>
<tr><td>1</td><td>SYNCI_Step</td><td>Address step size to be used with the SYNCI instruction. See that instruction’s description for the use of this value. In the typical implementation, this value should be zero if there are no caches in the system which must be synchronize (either because there are no caches, or because the instruction cache tracks writes to the data cache). In other cases, the return value should be the smallest line size of the caches that must be synchronize.</td><td>Required</td></tr>
<tr><td>2</td><td>CC</td><td>High-resolution cycle counter. This register provides read access to the coprocessor 0 Count Register.</td><td>Required</td></tr>
<tr><td>3</td><td>CCRes</td><td>Resolution of the CC register. This value denotes the number of cycles between update of the register. For example:
<table>
<tr><th>CCRes Value</th><th>Meaning</th></tr>
<tr><td>1</td><td>CC register increments every CPU cycle</td></tr>
<tr><td>2</td><td>CC register increments every second CPU cycle</td></tr>
<tr><td>3</td><td>CC register increments every third CPU cycle</td></tr>
<tr><td colspan="2">etc.</td></tr>
</table></td><td>Required</td></tr>
<tr><td>4-28</td><td></td><td>These registers numbers are reserved for future architecture use. Access results in a Reserved Instruction Exception.</td><td>Reserved</td></tr>
<tr><td>29</td><td>ULR</td><td>User Local Register. This register provides read access to the coprocessor 0 UserLocal register, if it is implemented. In some operating environments, the UserLocal register is a pointer to a thread-specific storage block.</td><td>Required if the UserLocal register is implemented</td></tr>
<tr><td>30-31</td><td></td><td>These register numbers are reserved for implementation-dependent use. If they are not implemented, access results in a Reserved Instruction Exception.</td><td>Optional</td></tr>
</table>

Using the *HWREna* register, privileged software may select which of the hardware registers are accessible via 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.

Software may determine which registers are implemented by writing all ones to the *HWREna* register, then reading the value back. If a bit reads back as a one, the processor implements that hardware register.

9.15 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/XTLB 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 9-17 shows the format of the *BadVAddr* register; Table 9.22 describes the *BadVAddr* register fields.

Figure 9-17 BadVAddr Register Format

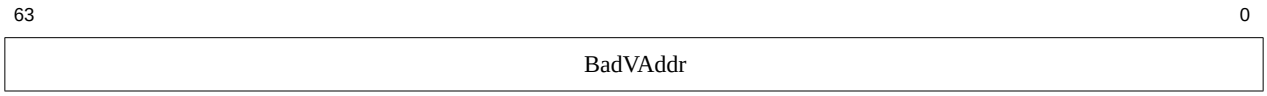

Table 9.22 BadVAddr Register Field Descriptions

Fields		Description	Read/Write	Reset State	Compliance
Name	Bits				
BadVAddr	63..0	Bad virtual address	R	Undefined	Required

9.16 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.

The Count register can also be read via RDHWR register 2.

Figure 9-18 shows the format of the *Count* register; Table 9.23 describes the *Count* register fields.

Figure 9-18 Count Register Format

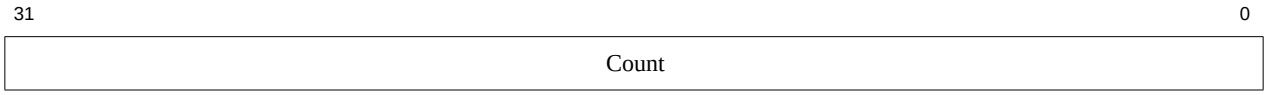

Table 9.23 Count Register Field Descriptions

Fields		Description	Read/Write	Reset State	Compliance
Name	Bits				
Count	31..0	Interval counter	R/W	Undefined	Required

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

Compliance Level: *Implementation Dependent*.

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

9.18 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, XTLB Refill, TLB Invalid, or TLB Modified) causes the bits of the virtual address corresponding to the *R* and *VPN2* fields to be written into 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.

Software may determine the value of *SEGBITS* by writing all ones to the *EntryHi* register and reading the value back. Bits read as "1" from the *VPN2* field allow software to determine the boundary between the *VPN2* and *Fill* fields to calculate the value of *SEGBITS*.

The *VPNX2*, *VPN2*, and *R* 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 or DMTC0) do not cause the implicit write of address-related fields in the *BadVAddr*, *Context*, or *XContext* registers.

Figure 9-19 shows the format of the *EntryHi* register; Table 9.24 describes the *EntryHi* register fields.

Figure 9-19 EntryHi Register Format

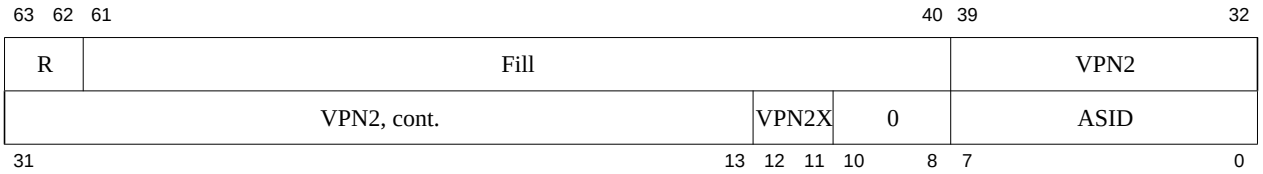

Table 9.24 EntryHi Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
R	63..62	Virtual memory region, corresponding to $VA_{63..62}$. **Encoding** / **Meaning** 0b00: xuseg: user address region 0b01: xsseg: supervisor address region. If Supervisor Mode is not implemented, this encoding is reserved 0b10: Reserved 0b11: xkseg: kernel address region This field is written by hardware on a TLB exception or on a TLB read, and is written by software before a TLB write. For processors implementing $Config_{AT}$ = 1 (access to 32-bit compatibility segments only), only the 0b00 and 0b11 values are legal. In this circumstance, the operation of the processor is **UNDEFINED** if $EntryHi_R$ is written with any other value, and the processor will only supply the legal values on an exception.	R/W	Undefined	Required
Fill	61..40	Fill bits reserved for expansion of the virtual address space. See below. Returns zeros on read, ignored on write.	R	0	Required
VPN2	39..13	$VA_{39..13}$ 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. The default width of this field implicitly limits the size of each virtual address space to 40 bits. If the processor implements fewer virtual address bits than this default, the Fill field must be extended to take up the unimplemented VPN2 bits. If the processor implements more virtual address bits than this default, the VPN2 field must be extended to take up some or all of the Fill bits.	R/W	Undefined	Required
VPN2X	12..11	In Release 2 of the Architecture (and subsequent releases), 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 $VA_{12..11}$ 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	10..8	Must be written as zero; returns zero on read.	0	0	Reserved

Table 9.24 EntryHi Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
ASID	7..0	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)

Programming Note:

In implementations of Release 2 (and subsequent releases) 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 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.

9.19 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 made. In Release 1 of the architecture, this request is combined in an implementation-dependent way with hardware interrupt 5 to set interrupt bit IP(7) in the *Cause* register. In Release 2 (and subsequent releases) of the Architecture, the presence of the interrupt is visible to software via the $Cause_{TI}$ bit and is combined in an implementation-dependent way with a hardware or software interrupt. For Vectored Interrupt Mode, the interrupt is at the level specified by the $IntCtl_{IPTI}$ field.

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 9-20 shows the format of the *Compare* register; Table 9.25 describes the *Compare* register fields.

Figure 9-20 Compare Register Format

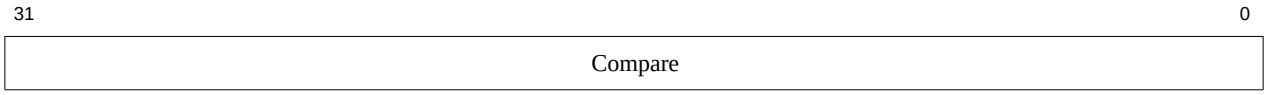

Table 9.25 Compare Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Compare	31..0	Interval count compare value	R/W	Undefined	Required

Programming Note:

In Release 2 of the Architecture, the EHB instruction can be used to make interrupt state changes visible when the *Compare* register is written. See 6.1.2.1 “Software Hazards and the Interrupt System” on page 58.

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

Compliance Level: *Implementation Dependent*.

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

9.21 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 "MIPS64 and microMIPS64 Operating Modes" on page 19 for a discussion of operating modes, and "Interrupts" on page 47 for a discussion of interrupt modes.

Figure 9-21 shows the format of the *Status* register; Table 9.26 describes the *Status* register fields.

Figure 9-21 Status Register Format

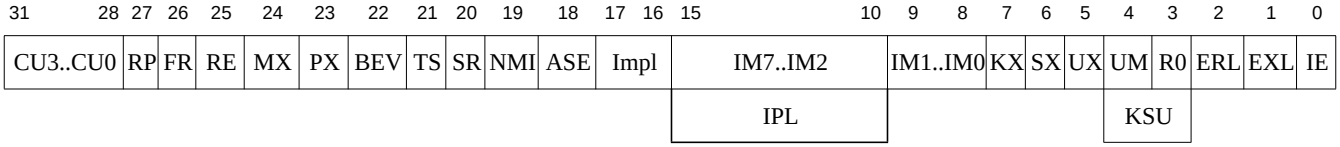

Table 9.26 Status Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
CU (CU3.. CU0)	31..28	Controls access to coprocessors 3, 2, 1, and 0, respectively: Encoding / Meaning: 0 – Access not allowed 1 – Access allowed Coprocessor 0 is always usable when the processor is running in Kernel Mode or Debug Mode, independent of the state of the CU_0 bit. In Release 2 (and subsequent releases) of the Architecture, and for 64-bit implementations of Release 1 of the Architecture, execution of all floating point instructions, including those encoded with the COP1X opcode, is controlled by the CU1 enable. CU3 is no longer used and is reserved for future use by the Architecture. If there is no provision for connecting a coprocessor, the corresponding CU bit must be ignored on write and read as zero.	R/W	Undefined	Required for all implemented coprocessors
RP	27	Enables reduced power mode on some implementations. The specific operation of this bit is implementation dependent. If this bit is not implemented, it must be ignored on write and read as zero. If this bit is implemented, the reset state must be zero so that the processor starts at full performance.	R/W	0	Optional

Table 9.26 Status Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>FR</td><td>26</td><td>In Release 1 of the Architecture, only MIPS64 processors could implement a 64-bit floating point unit. In Release 2 of the Architecture (and subsequent releases) , both 32-bit and 64-bit 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:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Floating point registers can contain any 32-bit datatype. 64-bit datatypes are stored in even-odd pairs of registers.</td></tr><tr><td>1</td><td>Floating point registers can contain any datatype</td></tr></table>This bit must be ignored on write and read as zero under the following conditions:
• 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 (and subsequent releases) in which a 64-bit floating point unit is not implemented
Certain combinations of the FR bit and other state or operations can cause UNPREDICTABLE behavior. See “64-bit FPR Enable” on page 21 for a discussion of these combinations.
When software changes the value of this bit, the contents of the floating point registers are UNPREDICTABLE.</td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>RE</td><td>25</td><td>Used to enable reverse-endian memory references while the processor is running in user mode:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>User mode uses configured endianness</td></tr><tr><td>1</td><td>User mode uses reversed endianness</td></tr></table>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.</td><td>R/W</td><td>Undefined</td><td>Optional</td></tr>
</table>

Table 9.26 Status Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>MX</td><td>24</td><td>Enables access to MDMX™ and MIPS® DSP resources on processors implementing one of these ASEs. If neither the MDMX nor the MIPS DSP ASE is implemented, this bit must be ignored on write and read as zero.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Access not allowed</td></tr><tr><td>1</td><td>Access allowed</td></tr></table></td><td>R if the processor implements neither the MDMX nor the MIPS DSP ASEs; otherwise R/W</td><td>0 if the processor implements neither the MDMX nor the MIPS DSP ASEs; otherwise Undefined</td><td>Optional</td></tr>
<tr><td>PX</td><td>23</td><td>Enables access to 64-bit operations in User mode, without enabling 64-bit addressing:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>64-bit operations are not enabled in User Mode</td></tr><tr><td>1</td><td>64-bit operations are enabled in User Mode</td></tr></table></td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>BEV</td><td>22</td><td>Controls the location of exception vectors:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Normal</td></tr><tr><td>1</td><td>Bootstrap</td></tr></table>See “Exception Vector Locations” on page 61 for details.</td><td>R/W</td><td>1</td><td>Required</td></tr>
</table>

Table 9.26 Status Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>TS[1]</td><td>21</td><td>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 (and subsequent releases), 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 “TLB Initialization” on page 35 for a discussion of software TLB initialization used to avoid a machine check exception 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 UNPREDICTABLE whether hardware ignores the write, accepts the write with no side effects, or accepts the write and initiates a machine check exception.</td><td>R/W</td><td>0</td><td>Required if the processor detects and reports a match on multiple TLB entries</td></tr>
<tr><td>SR</td><td>20</td><td>Indicates that the entry through the reset exception vector was due to a Soft Reset:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Not Soft Reset (NMI or Reset)</td></tr><tr><td>1</td><td>Soft Reset</td></tr></table>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.</td><td>R/W</td><td>1 for Soft Reset; 0 otherwise</td><td>Required if Soft Reset is implemented</td></tr>
<tr><td>NMI</td><td>19</td><td>Indicates that the entry through the reset exception vector was due to an NMI exception:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Not NMI (Soft Reset or Reset)</td></tr><tr><td>1</td><td>NMI</td></tr></table>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.</td><td>R/W</td><td>1 for NMI; 0 otherwise</td><td>Required if NMI is implemented</td></tr>
</table>

Table 9.26 Status Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
ASE	18	This bit is reserved for the MCU ASE. If MCU ASE is not implemented, then this bit must be written as zero; returns zero on read.	0 if MCU ASE is not imple-mented	0 if MCU ASE is not imple-mented	Required for MCU ASE; Otherwise Reserved
Impl	17..16	These bits are implementation dependent and are not defined by the architecture. If they are not implemented, they must be ignored on write and read as zero.		Undefined	Optional
IM7..IM2	15..10	Interrupt Mask: Controls the enabling of each of the hardware interrupts. Refer to "Interrupts" on page 47 for a complete discussion of enabled interrupts. Encoding / Meaning: 0 — Interrupt request disabled 1 — Interrupt request enabled In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled ($Config3_{VEIC}$ = 1), these bits take on a different meaning and are interpreted as the IPL field, described below.	R/W	Undefined	Required
IPL	15..10	Interrupt Priority Level. In implementations of Release 2 of the Architecture (and subsequent releases) in which EIC interrupt mode is enabled ($Config3_{VEIC}$ = 1), this field is the encoded (0..63) value of the current IPL. An interrupt will be signaled only if the requested IPL is higher than this value. If EIC interrupt mode is not enabled ($Config3_{VEIC}$ = 0), these bits take on a different meaning and are interpreted as the IM7..IM2 bits, described above.	R/W	Undefined	Optional (Release 2 and EIC interrupt mode only)
IM1..IM0	9..8	Interrupt Mask: Controls the enabling of each of the software interrupts. Refer to "Interrupts" on page 47 for a complete discussion of enabled interrupts. Encoding / Meaning: 0 — Interrupt request disabled 1 — Interrupt request enabled In implementations of Release 2 of the Architecture in which EIC interrupt mode is enabled ($Config3_{VEIC}$ = 1), these bits are writable, but have no effect on the interrupt system.	R/W	Undefined	Required

Table 9.26 Status Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>KX</td><td>7</td><td>Enables the following behavior:
• Access to 64-bit Kernel Segments
• Use of the XTLB Refill Vector for references to Kernel Segments
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Access to 64-bit Kernel Segments is disabled; TLB Refill Vector is used for references to Kernel Segments</td></tr><tr><td>1</td><td>Access to 64-bit Kernel Segments is enabled; XTLB Refill Vector is used for references to Kernel Segments</td></tr></table>If 64-bit addressing is not implemented, this bit must be ignored on write and read as zero.</td><td>R/W</td><td>Undefined</td><td>Required for 64-bit Addressing</td></tr>
<tr><td>SX</td><td>6</td><td>If Supervisor Mode is implemented, enables the following behavior:
• Access to 64-bit Supervisor Segments
• Use of the XTLB Refill Vector for references to Supervisor Segments
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Access to 64-bit Supervisor Segments is disabled; TLB Refill Vector is used for references to Supervisor Segments</td></tr><tr><td>1</td><td>Access to 64-bit Supervisor Segments is enabled; XTLB Refill Vector is used for references to Supervisor Segments</td></tr></table>If Supervisor Mode is not implemented, it is implementation dependent whether access to what would normally be 64-bit supervisor address space is enabled with the SX or KX bit.
If 64-bit addressing is not implemented, this bit must be ignored on write and read as zero.</td><td>R/W</td><td>Undefined</td><td>Required if both Supervisor Mode and 64-bit addressing are implemented</td></tr>
</table>

Table 9.26 Status Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>UX</td><td>5</td><td>Enables the following behavior:
• Access to 64-bit User Segments
• Use of the XTLB Refill Vector for references to User Segments
• Execution of instructions which perform 64-bit operations while the processor is operating in User Mode
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Access to 64-bit User Segments is disabled; TLB Refill Vector is used for references to User Segments; Execution of instructions which perform 64-bit operations is disallowed while the processor is running in User Mode</td></tr><tr><td>1</td><td>Access to 64-bit User Segments is enabled; XTLB Refill Vector is used for references to User Segments; Execution of instructions which perform 64-bit operations is allowed while the processor is running in User Mode</td></tr></table>If 64-bit addressing is not implemented, this bit must be ignored on write and read as zero.</td><td>R/W</td><td>Undefined</td><td>Required for 64-bit Addressing</td></tr>
<tr><td>KSU</td><td>4..3</td><td>If Supervisor Mode is implemented, the encoding of this field denotes the base operating mode of the processor. See “MIPS64 and microMIPS64 Operating Modes” on page 19 for a full discussion of operating modes. The encoding of this field is:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0b00</td><td>Base mode is Kernel Mode</td></tr><tr><td>0b01</td><td>Base mode is Supervisor Mode</td></tr><tr><td>0b10</td><td>Base mode is User Mode</td></tr><tr><td>0b11</td><td>Reserved. The operation of the processor is UNDEFINED if this value is written to the KSU field</td></tr></table>Note: This field overlaps the UM and R0 fields, described below.</td><td>R/W</td><td>Undefined</td><td>Required if Supervisor Mode is implemented; Optional otherwise</td></tr>
</table>

Table 9.26 Status Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>UM</td><td>4</td><td>If Supervisor Mode is not implemented, this bit denotes the base operating mode of the processor. See “MIPS64 and microMIPS64 Operating Modes” on page 19 for a full discussion of operating modes. The encoding of this bit is:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Base mode is Kernel Mode</td></tr><tr><td>1</td><td>Base mode is User Mode</td></tr></table>Note: This bit overlaps the KSU field, described above.</td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>R0</td><td>3</td><td>If Supervisor Mode is not implemented, this bit is reserved. This bit must be ignored on write and read as zero.
Note: This bit overlaps the KSU field, described above.</td><td>R</td><td>0</td><td>Reserved</td></tr>
<tr><td>ERL</td><td>2</td><td>Error Level; Set by the processor when a Reset, Soft Reset, NMI or Cache Error exception are taken.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Normal level</td></tr><tr><td>1</td><td>Error level</td></tr></table>When ERL is set:
• The processor is running in kernel mode
• Hardware and software interrupts are disabled
• The ERET instruction will use the return address held in ErrorEPC instead of EPC
• Segment kuseg is treated as an unmapped and uncached region. See “Address Translation for the kuseg Segment when StatusERL = 1” on page 33. 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.</td><td>R/W</td><td>1</td><td>Required</td></tr>
</table>

Table 9.26 Status Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
EXL	1	Exception Level; Set by the processor when any exception other than Reset, Soft Reset, NMI or Cache Error exception are taken. Encoding: 0 = Normal level; 1 = Exception level When EXL is set: • The processor is running in Kernel Mode • Hardware and software interrupts are disabled. • TLB/XTLB Refill exceptions use the general exception vector instead of the TLB/XTLB Refill vectors. • EPC, $Cause_{BD}$ and SRSCtl (implementations of Release 2 of the Architecture only) will not be updated if another exception is taken	R/W	Undefined	Required
IE	0	Interrupt Enable: Acts as the master enable for software and hardware interrupts: Encoding: 0 = Interrupts are disabled; 1 = Interrupts are enabled In Release 2 of the Architecture (and subsequent releases), this bit may be modified separately via the DI and EI instructions.	R/W	Undefined	Required

1. The TS bit originally indicated a “TLB Shutdown” condition in which circuits detected multiple TLB matches and shutdown the TLB to prevent physical damage. In newer designs, multiple TLB matches do not cause physical damage to the TLB structure, so the TS bit retains its name, but is simply an indicator to the machine check exception handler that multiple TLB matches were detected and reported by the processor.

Programming Note:

In Release 2 of the Architecture, the EHB instruction can be used to make interrupt state changes visible when the *IM*, *IPL*, *ERL*, *EXL*, or *IE* fields of the *Status* register are written. See “Software Hazards and the Interrupt System” on page 58.

9.22 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 9-22 shows the format of the *IntCtl* register; Table 9.27 describes the *IntCtl* register fields.

.

Figure 9-22 IntCtl Register Format

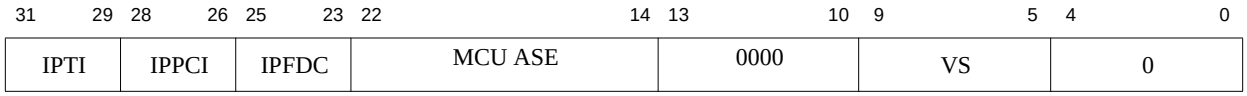

Table 9.27 IntCtl Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
IPTI	31..29	For Interrupt Compatibility and Vectored Interrupt modes, this field specifies the IP number to which the Timer Interrupt request is merged, and allows software to determine whether to consider $Cause_{TI}$ for a potential interrupt. Encoding / IP bit / Hardware Interrupt Source: 2 / 2 / HW0 3 / 3 / HW1 4 / 4 / HW2 5 / 5 / HW3 6 / 6 / HW4 7 / 7 / HW5 The value of this field is **UNPREDICTABLE** if External Interrupt Controller Mode is both implemented and enabled. The external interrupt controller is expected to provide this information for that interrupt mode.	R	Preset or Externally Set	Required

Table 9.27 IntCtl Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>IPPCI</td><td>28..26</td><td>For Interrupt Compatibility and Vectored Interrupt modes, this field specifies the IP number to which the Performance Counter Interrupt request is merged, and allows software to determine whether to consider $Cause_{PCI}$ for a potential interrupt.
<table><tr><th>Encoding</th><th>IP bit</th><th>Hardware Interrupt Source</th></tr><tr><td>2</td><td>2</td><td>HW0</td></tr><tr><td>3</td><td>3</td><td>HW1</td></tr><tr><td>4</td><td>4</td><td>HW2</td></tr><tr><td>5</td><td>5</td><td>HW3</td></tr><tr><td>6</td><td>6</td><td>HW4</td></tr><tr><td>7</td><td>7</td><td>HW5</td></tr></table>The value of this field is UNPREDICTABLE if External Interrupt Controller Mode is both implemented and enabled. The external interrupt controller is expected to provide this information for that interrupt mode.
If performance counters are not implemented ($Config1_{PC}$ = 0), this field returns zero on read.</td><td>R</td><td>Preset or Externally Set</td><td>Optional (Performance Counters Implemented)</td></tr>
<tr><td>IPFDC</td><td>25..23</td><td>For Interrupt Compatibility and Vectored Interrupt modes, this field specifies the IP number to which the Fast Debug Channel Interrupt request is merged, and allows software to determine whether to consider $Cause_{FDC}$ for a potential interrupt.
<table><tr><th>Encoding</th><th>IP bit</th><th>Hardware Interrupt Source</th></tr><tr><td>2</td><td>2</td><td>HW0</td></tr><tr><td>3</td><td>3</td><td>HW1</td></tr><tr><td>4</td><td>4</td><td>HW2</td></tr><tr><td>5</td><td>5</td><td>HW3</td></tr><tr><td>6</td><td>6</td><td>HW4</td></tr><tr><td>7</td><td>7</td><td>HW5</td></tr></table>The value of this field is UNPREDICTABLE if External Interrupt Controller Mode is both implemented and enabled. The external interrupt controller is expected to provide this information for that interrupt mode.
If EJTAG FDC is not implemented, this field returns zero on read.</td><td>R</td><td>Preset or Externally Set</td><td>Optional (EJTAG Fast Debug Channel Implemented)</td></tr>
</table>

Table 9.27 IntCtl Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>MCU ASE</td><td>22..14</td><td>These bits are reserved for the MicroController ASE.

If that ASE is not implemented, must be written as zero; returns zero on read.</td><td>0</td><td>0</td><td>Reserved</td></tr>
<tr><td>0</td><td>22..10</td><td>Must be written as zero; returns zero on read.</td><td>0</td><td>0</td><td>Reserved</td></tr>
<tr><td>VS</td><td>9..5</td><td>Vector Spacing. If vectored interrupts are implemented (as denoted by $Config3_{VInt}$ or $Config3_{VEIC}$), this field specifies the spacing between vectored interrupts.
<table>
<tr><th rowspan="2">Encoding</th><th colspan="2">Spacing Between Vectors</th></tr>
<tr><th>(hex)</th><th>(decimal)</th></tr>
<tr><td>0x00</td><td>0x000</td><td>0</td></tr>
<tr><td>0x01</td><td>0x020</td><td>32</td></tr>
<tr><td>0x02</td><td>0x040</td><td>64</td></tr>
<tr><td>0x04</td><td>0x080</td><td>128</td></tr>
<tr><td>0x08</td><td>0x100</td><td>256</td></tr>
<tr><td>0x10</td><td>0x200</td><td>512</td></tr>
</table>
All other values are reserved. The operation of the processor is UNDEFINED if a reserved value is written to this field.
If neither EIC interrupt mode nor VI mode are implemented ($Config3_{VEIC}$ = 0 and $Config3_{VINT}$ = 0), this field is ignored on write and reads as zero.</td><td>R/W</td><td>0</td><td>Optional</td></tr>
<tr><td>0</td><td>4..0</td><td>Must be written as zero; returns zero on read.</td><td>0</td><td>0</td><td>Reserved</td></tr>
</table>

9.23 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 9-23 shows the format of the *SRSCtl* register; Table 9.28 describes the *SRSCtl* register fields.

Figure 9-23 SRSCtl Register Format

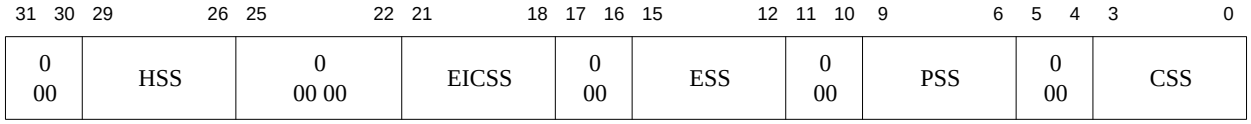

Table 9.28 SRSCtl Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
0	31..30	Must be written as zeros; returns zero on read.	0	0	Reserved
HSS	29..26	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. A non-zero value in this field indicates that the implemented shadow sets are numbered 0..n, where n is the value of the field. 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 *SRSMap* register. The operation of the processor is **UNDEFINED** if a value larger than the one in this field is written to any of these other values.	R	Preset	Required
0	25..22	Must be written as zeros; returns zero on read.	0	0	Reserved
EICSS	21..18	EIC interrupt mode shadow set. If $Config3_{VEIC}$ 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 *SRSMap* register to select the current shadow set for the interrupt. See "External Interrupt Controller Mode" on page 54 for a discussion of EIC interrupt mode. If $Config3_{VEIC}$ is 0, this field must be written as zero, and returns zero on read.	R	Undefined	Required (EIC interrupt mode only)
0	17..16	Must be written as zeros; returns zero on read.	0	0	Reserved

Table 9.28 SRSCtl Register Field Descriptions (Continued)

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
ESS	15..12	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 **UNDEFINED** if software writes a value into this field that is greater than the value in the HSS field.	R/W	0	Required
0	11..10	Must be written as zeros; returns zero on read.	0	0	Reserved
PSS	9..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 $Status_{BEV}$ = 0. This field is not updated on any exception which sets $Status_{ERL}$ to 1 (i.e., NMI or cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with $Status_{EXL}$ = 1, or $Status_{BEV}$ = 1. The operation of the processor is **UNDEFINED** if software writes a value into this field that is greater than the value in the HSS field.	R/W	0	Required
0	5..4	Must be written as zeros; returns zero on read.	0	0	Reserved
CSS	3..0	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 9.29 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 $Status_{ERL}$ to 1 (i.e., NMI or cache error), an entry into EJTAG Debug mode, or any exception or interrupt that occurs with $Status_{EXL}$ = 1, or $Status_{BEV}$ = 1. Neither is it updated on an ERET with $Status_{ERL}$ = 1 or $Status_{BEV}$ = 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 9.29 Sources for new $SRSCtl_{CSS}$ on an Exception or Interrupt

Exception Type	Condition	$SRSCtl_{CSS}$ Source	Comment
Exception	All	$SRSCtl_{ESS}$	

Table 9.29 Sources for new $SRSCtl_{CSS}$ on an Exception or Interrupt

Exception Type	Condition	$SRSCtl_{CSS}$ Source	Comment
Non-Vectored Interrupt	$Cause_{IV}$ = 0	$SRSCtl_{ESS}$	Treat as exception
Vectored Interrupt	$Cause_{IV}$ = 1 and $Config3_{VEIC}$ = 0 and $Config3_{VInt}$ = 1	$SRSMap_{VectNum \times 4+3..VectNum \times 4}$	Source is internal map register
Vectored EIC Interrupt	$Cause_{IV}$ = 1 and $Config3_{VEIC}$ = 1	$SRSCtl_{EICSS}$	Source is external interrupt controller.

Programming Note:

A software change to the PSS field creates an instruction hazard between the write of the *SRSCtl* register and the use of a RDPGPR or WRPGPR instruction. This hazard must be cleared with a JR.HB or JALR.HB instruction as described in "Hazard Clearing Instructions and Events" on page 88. A hardware change to the PSS field as the result of interrupt or exception entry is automatically cleared for the execution of the first instruction in the interrupt or exception handler.

9.24 SRSMap Register (CP0 Register 12, Select 3)

Compliance Level: *Required* in Release 2 (and subsequent releases) 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 ($Cause_{IV}$ = 0 or $IntCtl_{VS}$ = 0). In such cases, the shadow set number comes from $SRSCtl_{ESS}$.

If $SRSCtl_{HSS}$ 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 $SRSCtl_{HSS}$.

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 9-24 shows the format of the *SRSMap* register; Table 9.30 describes the *SRSMap* register fields.

Figure 9-24 SRSMap Register Format

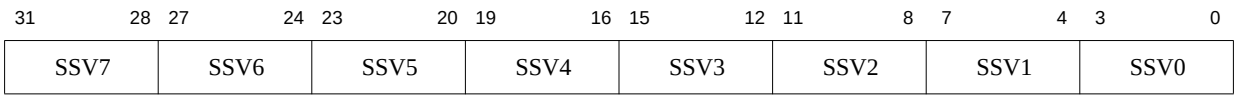

Table 9.30 SRSMap Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
SSV7	31..28	Shadow register set number for Vector Number 7	R/W	0	Required
SSV6	27..24	Shadow register set number for Vector Number 6	R/W	0	Required
SSV5	23..20	Shadow register set number for Vector Number 5	R/W	0	Required
SSV4	19..16	Shadow register set number for Vector Number 4	R/W	0	Required
SSV3	15..12	Shadow register set number for Vector Number 3	R/W	0	Required
SSV2	11..8	Shadow register set number for Vector Number 2	R/W	0	Required
SSV1	7..4	Shadow register set number for Vector Number 1	R/W	0	Required
SSV0	3..0	Shadow register set number for Vector Number 0	R/W	0	Required

9.25 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 9-25 shows the format of the *Cause* register; Table 9.31 describes the *Cause* register fields.

Figure 9-25 Cause Register Format

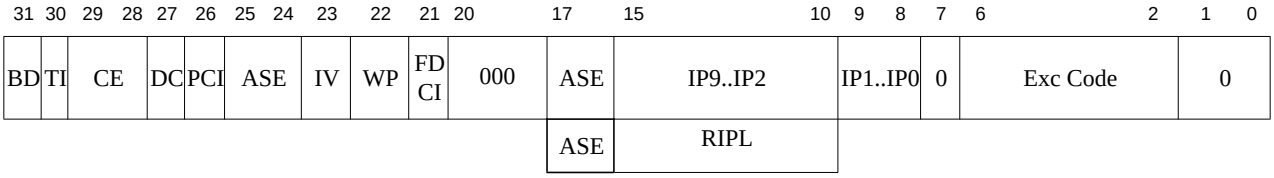

Table 9.31 Cause Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
BD	31	Indicates whether the last exception taken occurred in a branch delay slot: Encoding 0: Not in delay slot Encoding 1: In delay slot The processor updates BD only if $Status_{EXL}$ was zero when the exception occurred.	R	Undefined	Required
TI	30	Timer Interrupt. In an implementation of Release 2 of the Architecture, this bit denotes whether a timer interrupt is pending (analogous to the IP bits for other interrupt types): Encoding 0: No timer interrupt is pending Encoding 1: Timer interrupt is pending In an implementation of Release 1 of the Architecture, this bit must be written as zero and returns zero on read.	R	Undefined	Required (Release 2)

Table 9.31 Cause Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>CE</td><td>29..28</td><td>Coprocessor unit number referenced when a Coprocessor Unusable exception is taken. This field is loaded by hardware on every exception, but is UNPREDICTABLE for all exceptions except for Coprocessor Unusable.</td><td>R</td><td>Undefined</td><td>Required</td></tr>
<tr><td>DC</td><td>27</td><td>Disable Count register. In some power-sensitive applications, the Count register is not used but may still be the source of some noticeable power dissipation. This bit allows the Count register to be stopped in such situations.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Enable counting of Count register</td></tr><tr><td>1</td><td>Disable counting of Count register</td></tr></table>In an implementation of Release 1 of the Architecture, this bit must be written as zero, and returns zero on read.</td><td>R/W</td><td>0</td><td>Required (Release 2)</td></tr>
<tr><td>PCI</td><td>26</td><td>Performance Counter Interrupt. In an implementation of Release 2 of the Architecture (and subsequent releases), this bit denotes whether a performance counter interrupt is pending (analogous to the IP bits for other interrupt types):<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No performance counter interrupt is pending</td></tr><tr><td>1</td><td>Performance counter interrupt is pending</td></tr></table>In an implementation of Release 1 of the Architecture, or if performance counters are not implemented (Config1$_{PC}$ = 0), this bit must be written as zero and returns zero on read.</td><td>R</td><td>Undefined</td><td>Required (Release 2 and performance counters implemented)</td></tr>
<tr><td>ASE</td><td>25:24, 17:16</td><td>These bits are reserved for the MCU ASE.
If MCU ASE is not implemented, these bits return zero on reads and must be written with zeros.</td><td></td><td></td><td>Rrequired for MCU ASE; Otherwise Reserved</td></tr>
</table>

Table 9.31 Cause Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>IV</td><td>23</td><td>Indicates whether an interrupt exception uses the general exception vector or a special interrupt vector:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Use the general exception vector (0x180)</td></tr><tr><td>1</td><td>Use the special interrupt vector (0x200)</td></tr></table>In implementations of Release 2 of the architecture (and subsequent releases), if the $Cause_{IV}$ is 1 and $Status_{BEV}$ is 0, the special interrupt vector represents the base of the vectored interrupt table.</td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>WP</td><td>22</td><td>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.
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 the write, accepts the write with no side effects, or accepts the write and initiates a watch exception once $Status_{EXL}$ and $Status_{ERL}$ are both zero.
If watch registers are not implemented, this bit must be ignored on write and read as zero.</td><td>R/W</td><td>Undefined</td><td>Required if watch registers are implemented</td></tr>
<tr><td>FDCI</td><td>21</td><td>Fast Debug Channel Interrupt. This bit denotes whether a FDC interrupt is pending :
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No FDCinterrupt is pending</td></tr><tr><td>1</td><td>FDC interrupt is pending</td></tr></table></td><td>R</td><td>Undefined</td><td>Required</td></tr>
</table>

Table 9.31 Cause Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>IP7..IP2</td><td>15..10</td><td>Indicates an interrupt is pending:
<table><tr><th>Bit</th><th>Name</th><th>Meaning</th></tr><tr><td>15</td><td>IP7</td><td>Hardware interrupt 5</td></tr><tr><td>14</td><td>IP6</td><td>Hardware interrupt 4</td></tr><tr><td>13</td><td>IP5</td><td>Hardware interrupt 3</td></tr><tr><td>12</td><td>IP4</td><td>Hardware interrupt 2</td></tr><tr><td>11</td><td>IP3</td><td>Hardware interrupt 1</td></tr><tr><td>10</td><td>IP2</td><td>Hardware interrupt 0</td></tr></table>
In implementations of Release 1 of the Architecture, timer and performance counter interrupts are combined in an implementation-dependent way with hardware interrupt 5.
In implementations of Release 2 of the Architecture (and subsequent releases) in which EIC interrupt mode is not enabled ($Config3_{VEIC}$ = 0), timer and performance counter interrupts are combined in an implementation-dependent way with any hardware interrupt. If EIC interrupt mode is enabled ($Config3_{VEIC}$ = 1), these bits take on a different meaning and are interpreted as the RIPL field, described below.</td><td>R</td><td>Undefined</td><td>Required</td></tr>
<tr><td>RIPL</td><td>15..10</td><td>Requested Interrupt Priority Level.
In implementations of Release 2 of the Architecture (and subsequent releases) in which EIC interrupt mode is enabled ($Config3_{VEIC}$ = 1), this field is the encoded (0..63) value of the requested interrupt. A value of zero indicates that no interrupt is requested.
If EIC interrupt mode is not enabled ($Config3_{VEIC}$ = 0), these bits take on a different meaning and are interpreted as the IP7..IP2 bits, described above.</td><td>R</td><td>Undefined</td><td>Optional (Release 2 and EIC interrupt mode only)</td></tr>
<tr><td>IP1..IP0</td><td>9..8</td><td>Controls the request for software interrupts:
<table><tr><th>Bit</th><th>Name</th><th>Meaning</th></tr><tr><td>9</td><td>IP1</td><td>Request software interrupt 1</td></tr><tr><td>8</td><td>IP0</td><td>Request software interrupt 0</td></tr></table>
An implementation of Release 2 of the Architecture (and subsequent releases) which also implements EIC interrupt mode exports these bits to the external interrupt controller for prioritization with other interrupt sources.</td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>ExcCode</td><td>6..2</td><td>Exception code - see Table 9.32</td><td>R</td><td>Undefined</td><td>Required</td></tr>
<tr><td>0</td><td>25:24, 20..16, 7, 1..0</td><td>Must be written as zero; returns zero on read.</td><td>0</td><td>0</td><td>Reserved</td></tr>
</table>

Table 9.32 Cause Register ExcCode Field

Exception Code Value		Mnemonic	Description
Decimal	Hexadecimal		
0	0x00	Int	Interrupt
1	0x01	Mod	TLB modification exception
2	0x02	TLBL	TLB exception (load or instruction fetch)
3	0x03	TLBS	TLB exception (store)
4	0x04	AdEL	Address error exception (load or instruction fetch)
5	0x05	AdES	Address error exception (store)
6	0x06	IBE	Bus error exception (instruction fetch)
7	0x07	DBE	Bus error exception (data reference: load or store)
8	0x08	Sys	Syscall exception
9	0x09	Bp	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 $Debug_{DExcCode}$ field to denote an SDBBP in Debug Mode.
10	0x0a	RI	Reserved instruction exception
11	0x0b	CpU	Coprocessor Unusable exception
12	0x0c	Ov	Arithmetic Overflow exception
13	0x0d	Tr	Trap exception
14	0x0e	-	Reserved
15	0x0f	FPE	Floating point exception
16-17	0x10-0x11	-	Available for implementation dependent use
18	0x12	C2E	Reserved for precise Coprocessor 2 exceptions
19	0x13	TLBRI	TLB Read-Inhibit exception
20	0x14	TLBXI	TLB Execution-Inhibit exception
21	0x15	-	Reserved
22	0x16	MDMX	MDMX Unusable Exception (MDMX ASE)
23	0x17	WATCH	Reference to WatchHi/WatchLo address
24	0x18	MCheck	Machine check
25	0x19	Thread	Thread Allocation, Deallocation, or Scheduling Exceptions (MIPS® MT ASE)
26	0x1A	DSPDis	DSP ASE State Disabled exception (MIPS® DSP ASE)

Table 9.32 Cause Register ExcCode Field

Exception Code Value		Mnemonic	Description
Decimal	Hexadecimal		
27-29	0x20-0x1d	-	Reserved
30	0x1e	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 written to the $Debug_{DExcCode}$ field to indicate that re-entry to Debug Mode was caused by a cache error.
31	0x1f	-	Reserved

Programming Note:

In Release 2 of the Architecture (and the subsequent releases), the EHB instruction can be used to make interrupt state changes visible when the $IP_{1..0}$ field of the *Cause* register is written. See "Software Hazards and the Interrupt System" on page 58.

9.26 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.

Unless the *EXL* bit in the *Status* register is already a 1, the processor writes the *EPC* register when an exception occurs.

- 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 reads the *EPC* register as the result of execution of the ERET instruction.

Software may write the *EPC* register to change the processor resume address and read the *EPC* register to determine at what address the processor will resume.

Figure 9-26 shows the format of the *EPC* register; Table 9.33 describes the *EPC* register fields.

Figure 9-26 EPC Register Format

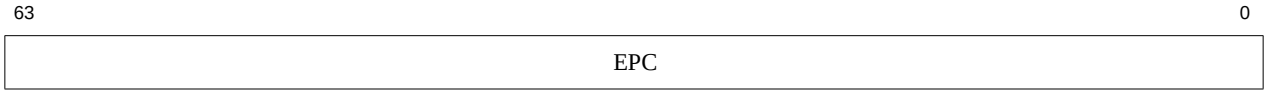

Table 9.33 EPC Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
EPC	63..0	Exception Program Counter	R/W	Undefined	Required

9.26.1 Special Handling of the EPC Register in Processors That Implement the MIPS16e ASE or the microMIPS64 Base Architectures

In processors that implement the MIPS16e ASE or microMIPS64 base architecture, the *EPC* register requires special handling.

When the processor writes the *EPC* register, it combines the address at which processing resumes with the value of the *ISA Mode* register:

```
EPC ← resumePC63..1 || ISAMode0
```

“resumePC” is the address at which processing resumes, as described above.

When the processor reads the *EPC* register, it distributes the bits to the *PC* and *ISAMode* registers:

```
PC ← EPC63..1 || 0
ISAMode ← EPC0
```

Software reads of the *EPC* register simply return to a GPR the last value written with no interpretation. Software writes to the *EPC* register store a new value which is interpreted by the processor as described above.

9.27 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 9-27 shows the format of the *PRId* register; Table 9.34 describes the *PRId* register fields.

Figure 9-27 PRId Register Format

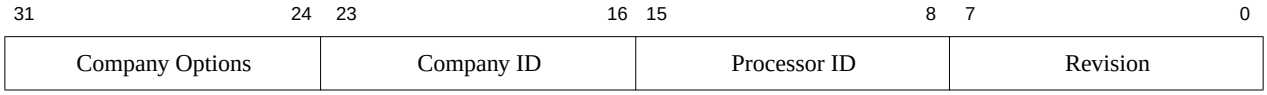

Table 9.34 PRId Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Company Options	31..24	Available to the designer or manufacturer of the processor for company-dependent options. The value in this field is not specified by the architecture. If this field is not implemented, it must read as zero.	R	Preset	Optional
Company ID	23..16	Identifies the company that designed or manufactured the processor. Software can distinguish a MIPS32/microMIPS32 or MIPS64/microMIPS64 processor from one implementing an earlier MIPS ISA by checking this field for zero. If it is non-zero the processor implements the MIPS32/microMIPS32 or MIPS64/microMIPS64 Architecture. Company IDs are assigned by MIPS Technologies when a MIPS32/microMIPS32 or MIPS64/microMIPS64 license is acquired. The encodings in this field are: **Encoding** / **Meaning** 0 — Not a MIPS32/microMIPS32 or MIPS64/microMIPS64 processor 1 — MIPS Technologies, Inc. 2-255 — Contact MIPS Technologies, Inc. for the list of Company ID assignments	R	Preset	Required
Processor ID	15..8	Identifies the type of processor. This field allows software to distinguish between various processor implementations within a single company, and is qualified by the CompanyID field, described above. The combination of the CompanyID and ProcessorID fields creates a unique number assigned to each processor implementation.	R	Preset	Required

Table 9.34 PRId Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Revision	7..0	Specifies the revision number of the processor. This field allows software to distinguish between one revision and another of the same processor type. If this field is not implemented, it must read as zero.	R	Preset	Optional

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 support@mips.com, reporting the specific case.

9.28 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 $Status_{BEV}$ is 0. The exception vector base address comes from the fixed defaults (see 6.2.2 "Exception Vector Locations" on page 61) when $Status_{BEV}$ 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 `0xFFFF.FFFF.8000.0000`, providing backward compatibility with Release 1 implementations.

Bits 31..30 of the *EBase* register are fixed with the value `0b10`, 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 9-28 shows the format of the *EBase* register; Table 9.35 describes the *EBase* register fields.

Figure 9-28 EBase Register Format

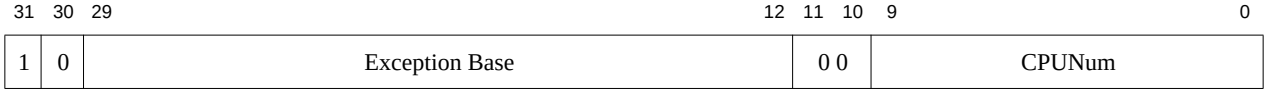

Table 9.35 EBase Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
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	29..12	In conjunction with bits 31..30, this field specifies the base address of the exception vectors when $Status_{BEV}$ is zero.	R/W	0	Required
0	11..10	Must be written as zero; returns zero on read.	0	0	Reserved

Table 9.35 EBase Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
CPUNum	9..0	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. This field can also be read via RDHWR register 0	R	Preset or Externally Set	Required

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 0xFFF 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 9.36 shows the conditions under which each *EBase* bit must be set to zero. *VN* represents the interrupt vector number as described in Table 6.4 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.

Table 9.36 Conditions Under Which EBase15..12 Must Be Zero

	Interrupt Vector Spacing in Bytes ($IntCtl_{VS}$[1])				
EBase bit	32	64	128	256	512
15	None	None	None	None	VN ≥ 63
14		None	None	VN ≥ 62	VN ≥ 31
13		None	VN ≥ 60	VN ≥ 30	VN ≥ 15
12		VN ≥ 56	VN ≥ 28	VN ≥ 14	VN ≥ 7

1. See Table 9.27 on page 144

9.29 CDMMBase Register (CP0 Register 15, Select 2)

Compliance Level: *Optional*.

The 64-bit physical base address for the Common Device Memory Map facility is defined by this register. This register only exists if $Config3_{CDMM}$ is set to one.

For devices that implement multiple VPEs, access to this register is controlled by the $VPEConf0_{MVP}$ register field. If the MVP bit is cleared, a read to this register returns all zeros and a write to this register is ignored.

Figure 9.29 has the format of the *CDMMBase* register, and Table 9.37 describes the register fields.

Figure 9.29 CDMMBase Register

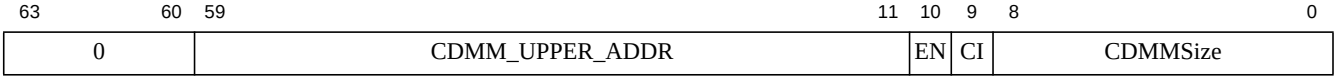

Table 9.37 CDMMBase Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
0	63:60	Must be written as zero; returns zero on read	0	0	Reserved
CDMM_UPPER_ADDR	59:11	Bits 63:15 of the base physical address of the memory mapped registers. The number of implemented physical address bits is implementation specific, see Section "Physical Memory" on page 23. For the unimplemented address bits - writes are ignored, returns zero on read.	R/W	Undefined	Required
EN	10	Enables the CDMM region. If this bit is cleared, memory requests to this address region go to regular system memory. If this bit is set, memory requests to this region go to the CDMM logic Encoding / Meaning: 0 — CDMM Region is disabled. 1 — CDMM Region is enabled.	R/W	0	Required
CI	9	If set to 1, this indicates that the first 64-byte Device Register Block of the CDMM is reserved for additional registers which manage CDMM region behavior and are not IO device registers.	R	Preset	Optional

Table 9.37 CDMMBase Register Field Descriptions (Continued)

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>CDMMSize</td><td>8:0</td><td>This field represents the number of 64-byte Device Register Blocks are instantiated in the core.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>1 DRB</td></tr><tr><td>1</td><td>2 DRBs</td></tr><tr><td>2</td><td>3 DRBs</td></tr><tr><td>...</td><td>...</td></tr><tr><td>511</td><td>512 DRBs</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

9.30 CMGCRBase Register (CP0 Register 15, Select 3)

Compliance Level: *Optional*.

The 64-bit physical base address for the memory-mapped Coherency Manager Global Configuration Register space is reflected by this register. This register only exists if *Config3*$_{CMGCR}$ is set to one.

On devices that implement the MIPS MT ASE, this register is instantiated once per processor.

Figure 9.30 has the format of the *CMGCRBase* register, and Table 9.38 describes the register fields.

Figure 9.30 CMGCRBase Register

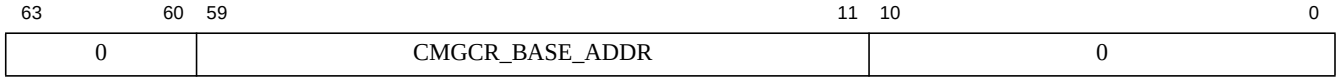

Table 9.38 CMGCRBase Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
CMGCR_BASE_ADDR	59:11	Bits 63:15 of the base physical address of the memory-mapped Coherency Manager GCR registers. This register field reflects the value of the GCR_BASE field within the memory-mapped Coherency Manager GCR Base Register. The number of implemented physical address bits is implementation specific, see Section "Physical Memory" on page 23. For the unimplemented address bits - writes are ignored, returns zero on read.	R	Preset (IP Configuration Value)	Required
0	63:60, 10:0	Must be written as zero; returns zero on read	0	0	Reserved

9.31 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. Three fields, *K23*, *KU*, and *K0*, must be initialized by software in the reset exception handler.

Figure 9-31 shows the format of the *Config* register; Table 9.39 describes the *Config* register fields.

Figure 9-31 Config Register Format

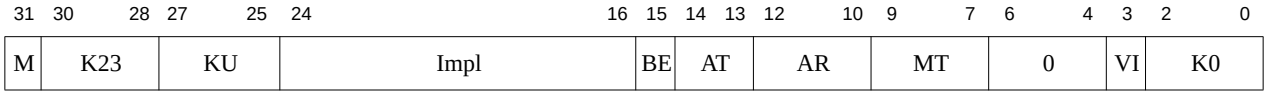

Table 9.39 Config Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
M	31	Denotes that the *Config1* register is implemented at a select field value of 1.	R	1	Required
K23	30:28	For processors that implement a Fixed Mapping MMU, this field specifies the kseg2 and kseg3 cacheability and coherency attribute. For processors that do not implement a Fixed Mapping MMU, this field reads as zero and is ignored on write. See "Alternative MMU Organizations" on page 215 for a description of the Fixed Mapping MMU organization.	R/W	Undefined for processors with a Fixed Mapping MMU; 0 otherwise	Optional
KU	27:25	For processors that implement a Fixed Mapping MMU, this field specifies the kuseg cacheability and coherency attribute. For processors that do not implement a Fixed Mapping MMU, this field reads as zero and is ignored on write. See "Alternative MMU Organizations" on page 215 for a description of the Fixed Mapping MMU organization.	R/W	Undefined for processors with a Fixed Mapping MMU; 0 otherwise	Optional
Impl	24:16	This field is reserved for implementations. Refer to the processor specification for the format and definition of this field		Undefined	Optional
BE	15	Indicates the endian mode in which the processor is running: Encoding: 0 — Little endian; 1 — Big endian	R	Preset or Externally Set	Required

Table 9.39 Config Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>AT</td><td>14:13</td><td>Architecture Type implemented by the processor.

For Release 3, encoding values of 0-2, denotes address and register width (32-bit or 64-bit).

The implemented instruction sets (MIPS32/64 and/or microMIPS32/64) are denoted by the ISA register field of Config3.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>MIPS32 or microMIPS32</td></tr><tr><td>1</td><td>MIPS64 or microMIPS64 with access only to 32-bit compatibility segments</td></tr><tr><td>2</td><td>MIPS64or microMIPS64 with access to all address segments</td></tr><tr><td>3</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>AR</td><td>12:10</td><td>MIPS64 Architecture revision level.

microMIPS64 Architecture revision level is denoted by the MMAR field of Config3. If Config3 register is not implemented then microMIPS is not implemented.

If the ISA field of Config3 is one, then MIPS64 is not implemented and this field is not used.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Release 1</td></tr><tr><td>1</td><td>Release 2 or Release 3/MIPSr3
All features introduced in Release 3 are optional and detectable through Config3 register fields.</td></tr><tr><td>2-7</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

Table 9.39 Config Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>MT</td><td>9:7</td><td>MMU Type:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>None</td></tr><tr><td>1</td><td>Standard TLB (See “TLB Organization” on page 34)</td></tr><tr><td>2</td><td>BAT (See “Block Address Translation” on page 219)</td></tr><tr><td>3</td><td>Fixed Mapping (See “Fixed Mapping MMU” on page 215)</td></tr><tr><td>4</td><td>Dual VTLB and FTLB (See “Dual Variable-Page-Size and Fixed-Page-Size TLBs” on page 222)</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>0</td><td>6:4</td><td>Must be written as zero; returns zero on read.</td><td>0</td><td>0</td><td>Reserved</td></tr>
<tr><td>VI</td><td>3</td><td>Virtual instruction cache (using both virtual indexing and virtual tags):
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Instruction Cache is not virtual</td></tr><tr><td>1</td><td>Instruction Cache is virtual</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>K0</td><td>2:0</td><td>Kseg0 cacheability and coherency attribute. See Table 9.9 on page 106 for the encoding of this field.</td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
</table>

9.32 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 9-1 shows the format of the *Config1* register; Table 9-1 describes the *Config1* register fields.

Figure 9-1 Config1 Register Format

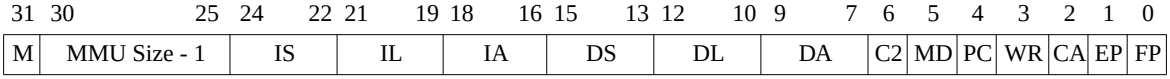

Table 9-1 Config1 Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	**Bits**				
M	31	This bit is reserved to indicate that a *Config2* register is present. If the *Config2* register is not implemented, this bit should read as a 0. If the *Config2* register is implemented, this bit should read as a 1.	R	Preset	Required
MMU Size - 1	30..25	Number of entries in the TLB minus one. The values 0 through 63 in this field correspond to 1 to 64 TLB entries. The value zero is implied by $Config_{MT}$ having a value of ‘none’.	R	Preset	Required
IS	24:22	Icache sets per way: **Encoding** / **Meaning** 0 – 64 1 – 128 2 – 256 3 – 512 4 – 1024 5 – 2048 6 – 4096 7 – 32	R	Preset	Required

Table 9-1 Config1 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read/ Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>IL</td><td>21:19</td><td>Icache line size:<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No Icache present</td></tr><tr><td>1</td><td>4 bytes</td></tr><tr><td>2</td><td>8 bytes</td></tr><tr><td>3</td><td>16 bytes</td></tr><tr><td>4</td><td>32 bytes</td></tr><tr><td>5</td><td>64 bytes</td></tr><tr><td>6</td><td>128 bytes</td></tr><tr><td>7</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>IA</td><td>18:16</td><td>Icache associativity:<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Direct mapped</td></tr><tr><td>1</td><td>2-way</td></tr><tr><td>2</td><td>3-way</td></tr><tr><td>3</td><td>4-way</td></tr><tr><td>4</td><td>5-way</td></tr><tr><td>5</td><td>6-way</td></tr><tr><td>6</td><td>7-way</td></tr><tr><td>7</td><td>8-way</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>DS</td><td>15:13</td><td>Dcache sets per way:<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>64</td></tr><tr><td>1</td><td>128</td></tr><tr><td>2</td><td>256</td></tr><tr><td>3</td><td>512</td></tr><tr><td>4</td><td>1024</td></tr><tr><td>5</td><td>2048</td></tr><tr><td>6</td><td>4096</td></tr><tr><td>7</td><td>32</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

Table 9-1 Config1 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read/ Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>DL</td><td>12:10</td><td>Dcache line size:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No Dcache present</td></tr><tr><td>1</td><td>4 bytes</td></tr><tr><td>2</td><td>8 bytes</td></tr><tr><td>3</td><td>16 bytes</td></tr><tr><td>4</td><td>32 bytes</td></tr><tr><td>5</td><td>64 bytes</td></tr><tr><td>6</td><td>128 bytes</td></tr><tr><td>7</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>DA</td><td>9:7</td><td>Dcache associativity:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Direct mapped</td></tr><tr><td>1</td><td>2-way</td></tr><tr><td>2</td><td>3-way</td></tr><tr><td>3</td><td>4-way</td></tr><tr><td>4</td><td>5-way</td></tr><tr><td>5</td><td>6-way</td></tr><tr><td>6</td><td>7-way</td></tr><tr><td>7</td><td>8-way</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>C2</td><td>6</td><td>Coprocessor 2 implemented:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No coprocessor 2 implemented</td></tr><tr><td>1</td><td>Coprocessor 2 implements</td></tr></table>
This bit indicates not only that the processor contains support for Coprocessor 2, but that such a coprocessor is attached.</td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>MD</td><td>5</td><td>MDMX ASE implemented:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No MDMX ASE implemented</td></tr><tr><td>1</td><td>MDMX ASE implemented</td></tr></table>
This bit indicates not only that the processor contains support for MDMX, but that such a processing element is attached.</td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

Table 9-1 Config1 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read/ Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>PC</td><td>4</td><td>Performance Counter registers implemented:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No performance counter registers implemented</td></tr><tr><td>1</td><td>Performance counter registers implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>WR</td><td>3</td><td>Watch registers implemented:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No watch registers implemented</td></tr><tr><td>1</td><td>Watch registers implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>CA</td><td>2</td><td>Code compression (MIPS16e) implemented:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>MIPS16e not implemented</td></tr><tr><td>1</td><td>MIPS16e implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>EP</td><td>1</td><td>EJTAG implemented:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No EJTAG implemented</td></tr><tr><td>1</td><td>EJTAG implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>FP</td><td>0</td><td>FPU implemented:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>No FPU implemented</td></tr><tr><td>1</td><td>FPU implemented</td></tr></table>
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 FIR CP1 register.</td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

9.33 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 9-32 shows the format of the *Config2* register; Table 9.40 describes the *Config2* register fields.

Figure 9-32 Config2 Register Format

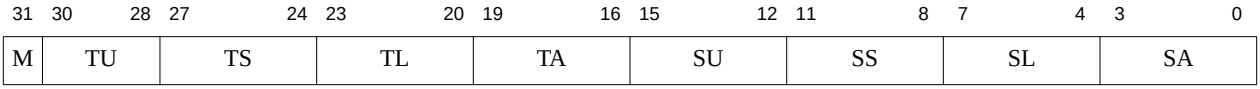

Table 9.40 Config2 Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
M	31	This bit is reserved to indicate that a *Config3* register is present. If the *Config3* register is not implemented, this bit should read as a 0. If the *Config3* register is implemented, this bit should read as a 1.	R	Preset	Required
TU	30:28	Implementation-specific tertiary cache control or status bits. If this field is not implemented it should read as zero and be ignored on write.	R/W	Preset	Optional
TS	27:24	Tertiary cache sets per way: 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 9.40 Config2 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>TL</td><td>23:20</td><td>Tertiary cache line size:<table><tr><th>Encoding</th><th>Line Size</th></tr><tr><td>0</td><td>No cache present</td></tr><tr><td>1</td><td>4</td></tr><tr><td>2</td><td>8</td></tr><tr><td>3</td><td>16</td></tr><tr><td>4</td><td>32</td></tr><tr><td>5</td><td>64</td></tr><tr><td>6</td><td>128</td></tr><tr><td>7</td><td>256</td></tr><tr><td>8-15</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>TA</td><td>19:16</td><td>Tertiary cache associativity:<table><tr><th>Encoding</th><th>Associativity</th></tr><tr><td>0</td><td>Direct Mapped</td></tr><tr><td>1</td><td>2</td></tr><tr><td>2</td><td>3</td></tr><tr><td>3</td><td>4</td></tr><tr><td>4</td><td>5</td></tr><tr><td>5</td><td>6</td></tr><tr><td>6</td><td>7</td></tr><tr><td>7</td><td>8</td></tr><tr><td>8-15</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>SU</td><td>15:12</td><td>Implementation-specific secondary cache control or status bits. If this field is not implemented it should read as zero and be ignored on write.</td><td>R/W</td><td>Preset</td><td>Optional</td></tr>
</table>

Table 9.40 Config2 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>SS</td><td>11:8</td><td>Secondary cache sets per way:<table><tr><th>Encoding</th><th>Sets Per Way</th></tr><tr><td>0</td><td>64</td></tr><tr><td>1</td><td>128</td></tr><tr><td>2</td><td>256</td></tr><tr><td>3</td><td>512</td></tr><tr><td>4</td><td>1024</td></tr><tr><td>5</td><td>2048</td></tr><tr><td>6</td><td>4096</td></tr><tr><td>7</td><td>8192</td></tr><tr><td>8-15</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>SL</td><td>7:4</td><td>Secondary cache line size:<table><tr><th>Encoding</th><th>Line Size</th></tr><tr><td>0</td><td>No cache present</td></tr><tr><td>1</td><td>4</td></tr><tr><td>2</td><td>8</td></tr><tr><td>3</td><td>16</td></tr><tr><td>4</td><td>32</td></tr><tr><td>5</td><td>64</td></tr><tr><td>6</td><td>128</td></tr><tr><td>7</td><td>256</td></tr><tr><td>8-15</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>SA</td><td>3:0</td><td>Secondary cache associativity:<table><tr><th>Encoding</th><th>Associativity</th></tr><tr><td>0</td><td>Direct Mapped</td></tr><tr><td>1</td><td>2</td></tr><tr><td>2</td><td>3</td></tr><tr><td>3</td><td>4</td></tr><tr><td>4</td><td>5</td></tr><tr><td>5</td><td>6</td></tr><tr><td>6</td><td>7</td></tr><tr><td>7</td><td>8</td></tr><tr><td>8-15</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

9.34 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 9-33 shows the format of the *Config3* register; Table 9.41 describes the *Config3* register fields.

Figure 9-33 Config3 Register Format

31	30	29	28 … 23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
M	BPG	CMGCR	0 0000000	IPLW		MMAR			MuCon	ISA On Exc	ISA		ULRI	RXI	DSP2P	DSPP	CTXTC	ITL	LPA	VEIC	VInt	SP	CDMM	MT	SM	TL

Table 9.41 Config3 Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
M	31	This bit is reserved to indicate that a *Config4* register is present. If the *Config4* register is not implemented, this bit should read as a 0. If the *Config4* register is implemented, this bit should read as a 1.	R	Preset	Required
BPG	30	Big Pages feature is implemented. This bit indicates that TLB pages larger than 256 MB are supported and that *C0_PageMask* Register is 64-bits wide. Encoding 0: Big Pages are not implemented and PageMask register is 32bits wide. Encoding 1: Big Pages are implemented and PageMask register is 64bits wide.	R	Preset	Required

Table 9.41 Config3 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>CMGCR</td><td>29</td><td>Coherency Manager memory-mapped Global Configuration Register Space is implemented.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>CM GCR space is not implemented</td></tr><tr><td>1</td><td>CM GCR space is implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required for Coherent Multiple-Core implementations that use the Coherency Manager.</td></tr>
<tr><td>0</td><td>28:23, 12, 9, 3</td><td>Must be written as zeros; returns zeros on read</td><td>0</td><td>0</td><td>Reserved</td></tr>
<tr><td>IPLW</td><td>22:21</td><td>Width of $Status_{IPL}$ and $Cause_{RIPL}$ fields:
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>IPL and RIPL fields are 6-bits in width.</td></tr><tr><td>1</td><td>IPL and RIPL fields are 8-bits in width.</td></tr><tr><td>Others</td><td>Reserved.</td></tr></table>If the IPL field is 8-bits in width, bits 18 and 16 of Status are used as the most significant bit and second most significant bit, respectively, of that field.

If the RIPL field is 8-bits in width, bits 17 and 16 of Cause are used as the most significant bit and second most significant bit, respectively, of that field.</td><td>R</td><td>Preset</td><td>Required if MCU ASE is implemented</td></tr>
<tr><td>MMAR</td><td>20:18</td><td>microMIPS64 Architecture revision level.

MIPS64 Architecture revision level is denoted by the AR field of Config.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Release3/MIPSr3</td></tr><tr><td>1-7</td><td>Reserved</td></tr></table>If ISA field of Config3 is zero, then microMIPS64 is not implemented and this field is not used.</td><td>R</td><td>Preset</td><td>Required if microMIPS is implemented</td></tr>
</table>

Table 9.41 Config3 Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
MCU	17	MIPS® MCU ASE is implemented. Encoding: 0 — MCU ASE is not implemented.; 1 — MCU ASE is implemented	R	Preset	Required if MCU ASE is implemented
ISAOnExc	16	Reflects the Instruction Set Architecture used after vectoring to an exception. Affects all exceptions whose offsets are relative to EBase. Encoding: 0 — MIPS64 is used on entrance to an exception vector.; 1 — microMIPS is used on entrance to an exception vector.	RW	Undefined	Required if microMIPS is implemented
ISA	15:14	Indicates Instruction Set Availability. Encoding: 0 — Only MIPS64 Instruction Set is implemented.; 1 — Only microMIPS64 is implemented.; 2 — Both MIPS64 and microMIPS64 ISAs are implemented. MIPS64 ISA used when coming out of reset.; 3 — Both MIPS64 and microMIPS64 ISAs are implemented. microMIPS64 ISA used when coming out of reset.	R	Preset	Required if microMIPS is implemented
ULRI	13	*UserLocal* register implemented. This bit indicates whether the *UserLocal* coprocessor 0 register is implemented. Encoding: 0 — *UserLocal* register is not implemented; 1 — *UserLocal* register is implemented	R	Preset	Required

Table 9.41 Config3 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>RXI</td><td>12</td><td>Indicates whether the RIE and XIE bits exist within the PageGrain register.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>The RIE and XIE bits are not implemented within the PageGrain register.</td></tr><tr><td>1</td><td>The RIE and XIE bits are implemented within the PageGrain register.</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>DSP2P</td><td>11</td><td>MIPS® DSP ASE Revision 2 implemented. This bit indicates whether Revision 2 of the MIPS DSP ASE is implemented.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Revision 2 of the MIPS DSP ASE is not implemented</td></tr><tr><td>1</td><td>Revision 2 of the MIPS DSP ASE is implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>DSPP</td><td>10</td><td>MIPS® DSP ASE implemented. This bit indicates whether the MIPS DSP ASE is implemented.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>MIPS DSP ASE is not implemented</td></tr><tr><td>1</td><td>MIPS DSP ASE is implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>CTXTC</td><td>9</td><td>ContextConfig and XContextConfig registers are implemented and the width of the BadVPN2 field within the Config register and the XConfig register depends on the contents of the ContextConfig register and XContextConfig register respectively.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>ContextConfig and XContextConfig are not implemented.</td></tr><tr><td>1</td><td>ContextConfig and XContextConfig are implemented and is used for the $Config_{BadVPN2}$ and $XConfig_{BadVPN2}$ fields.</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

Table 9.41 Config3 Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
ITL	8	MIPS® IFlowTrace™ mechanism implemented. This bit indicates whether the MIPS IFlowTrace is implemented. Encoding 0: MIPS IFlowTrace is not implemented Encoding 1: MIPS IFlowTrace is implemented	R	Preset	Required (Release 2.1 Only)
LPA	7	Large physical address support is implemented, and the *PageGrain* register exists Encoding 0: Large physical address support is not implemented Encoding 1: Large physical address support is implemented 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 0: Support for EIC interrupt mode is not implemented Encoding 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 0: Vector interrupts are not implemented Encoding 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 9.41 Config3 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>SP</td><td>4</td><td>Small (1KByte) page support is implemented, and the PageGrain register exists<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Small page support is not implemented</td></tr><tr><td>1</td><td>Small page support is implemented</td></tr></table>For implementations of Release 1 of the Architecture, this bit returns zero on read.</td><td>R</td><td>Preset</td><td>Required (Release 2 Only)</td></tr>
<tr><td>CDMM</td><td>3</td><td>Common Device Memory Map implemented. This bit indicates whether the CDMM is implemented.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>CDMM is not implemented</td></tr><tr><td>1</td><td>CDMM is implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>MT</td><td>2</td><td>MIPS® MT ASE implemented. This bit indicates whether the MIPS MT ASE is implemented.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>MIPS MT ASE is not implemented</td></tr><tr><td>1</td><td>MIPS MT ASE is implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>SM</td><td>1</td><td>SmartMIPS™ ASE implemented. This bit indicates whether the SmartMIPS ASE is implemented.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>SmartMIPS ASE is not implemented</td></tr><tr><td>1</td><td>SmartMIPS ASE is implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>TL</td><td>0</td><td>Trace Logic implemented. This bit indicates whether PC or data trace is implemented.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Trace logic is not implemented</td></tr><tr><td>1</td><td>Trace logic is implemented</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
</table>

9.35 Configuration Register 4 (CP0 Register 16, Select 4)

Compliance Level: *Required* if any optional feature described by this register is implemented: Release 2 of the Architecture; *Optional* otherwise.

The *Config4* register encodes additional capabilities.

The number of page-pair entries within the FTLB = decode(FTLBSets) * decode(FTLBWays).

Figure 9-34 shows the format of the *Config4* register; Table 9.42 describes the *Config4* register fields.

Figure 9-34 Config4 Register Format

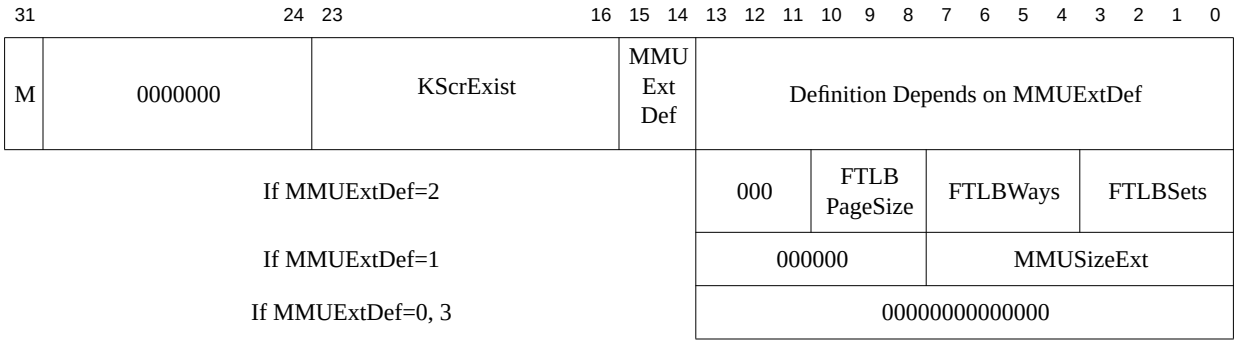

Table 9.42 Config4 Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
M	31	This bit is reserved to indicate that a *Config5* register is present. With the current architectural definition, this bit should always read as a 0.	R	Preset	Required
0	30:24	Must be written as zeros; returns zeros on read	R	0	Reserved
KScr Exist	23:16	Indicates how many scratch registers are available to kernel-mode software within COP0 Register 31. Each bit represents a select for Coproecessor0 Register 31. Bit 16 represents Select 0, Bit 23 represents Select 7. If the bit is set, the associated scratch register is implemented and available for kernel-mode software. Scratch registers meant for other purposes are not represented in this field. For example, if EJTAG is implemented, Bit 16 is preset to zero eventhough DESAVE register is implemented at Select 0. Select 1 is reserved for future debug purposes and should not be used as a kernel scratch register, so bit 17 is preset to zero.	R	Preset	Required if Kernel Scratch Registers are available

Table 9.42 Config4 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>MMU Ext Def</td><td>15:14</td><td>MMU Extension Definition.
Defines how Config4[13:0] is to be interpreted.

<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>1</td><td>Config4[7:0] used as MMUSizeExt.</td></tr><tr><td>2</td><td>Config4[3:0] used as FTLBSets.
Config4[7:4] used as FTLBWays.
Config4[10:8] used as FTLBPageSize.</td></tr><tr><td>0, 3</td><td>Reserved.
Config4[13:0] - Must be written as zeros, returns zeros on read.</td></tr></table></td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>FTLB Page Size</td><td>10:8</td><td>Indicates the Page Size of the FTLB Array Entries.
<table><tr><th>Encoding</th><th>Page Size</th></tr><tr><td>0</td><td>1 KB</td></tr><tr><td>1</td><td>4 KB</td></tr><tr><td>2</td><td>16 KB</td></tr><tr><td>3</td><td>64KB</td></tr><tr><td>4</td><td>256 KB</td></tr><tr><td>5</td><td>1 GB</td></tr><tr><td>6</td><td>4 GB</td></tr><tr><td>7</td><td>Reserved</td></tr></table>Implementations are allowed to implement any subset of these sizes, even a subset of only one pagesize. Software can detect if a FTLB page size is implemented by writing the desired size into this register field. If the size is implemented, the register field is updated to the desired encoding. If the size is not implemented, the register field value is not changed.

The FTLB must be flushed of any valid entries before this register field value is changed by software. The FTLB behavior is UNDEFINED if there are valid FTLB entries which were not all programmed using a common page size.</td><td>RW if multiple FTLB page-sizes are implle-mented

R if only one FTLB page size is imple-mented.</td><td>Preset, chosen value is implemen-tation spe-cific</td><td>Required if MMUExt-Def=2</td></tr>
</table>

Table 9.42 Config4 Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>FTLB Ways</td><td>7:4</td><td>Indicates the Set Associativity of the FTLB Array.<table><tr><th>Encoding</th><th>Associativity</th></tr><tr><td>0</td><td>2</td></tr><tr><td>1</td><td>3</td></tr><tr><td>2</td><td>4</td></tr><tr><td>3</td><td>5</td></tr><tr><td>4</td><td>6</td></tr><tr><td>5</td><td>7</td></tr><tr><td>6</td><td>8</td></tr><tr><td>7-15</td><td>Reserved</td></tr></table></td><td>R</td><td>Preset</td><td>Required if MMUExt-Def=2</td></tr>
<tr><td>FTLB Sets</td><td>3:0</td><td>Indicates the number of Sets per Way within the FTLB Array.<table><tr><th>Encoding</th><th>Sets per Way</th></tr><tr><td>0</td><td>1</td></tr><tr><td>1</td><td>2</td></tr><tr><td>2</td><td>4</td></tr><tr><td>3</td><td>8</td></tr><tr><td>4</td><td>16</td></tr><tr><td>5</td><td>32</td></tr><tr><td>6</td><td>64</td></tr><tr><td>7</td><td>128</td></tr><tr><td>8</td><td>256</td></tr><tr><td>9</td><td>512</td></tr><tr><td>10</td><td>1024</td></tr><tr><td>11</td><td>2048</td></tr><tr><td>12</td><td>4096</td></tr><tr><td>13</td><td>8192</td></tr><tr><td>14</td><td>16384</td></tr><tr><td>15</td><td>32768</td></tr></table></td><td>R</td><td>Preset</td><td>Required if MMUExt-Def=2</td></tr>
</table>

Table 9.42 Config4 Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
MMU Size Ext	7:0	If $Config4_{MMUExt}$=1 then this field is an extension of $Config1_{MMUSize-1}$ field. This field is concatenated to the left of the most significant bit to the MMUSize-1 field to indicate the size of the TLB-1.	R	Preset	Required if MMUExt-Def=1

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

Compliance Level: *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.

The architecture only defines the use of the M bits for presence detection of Selects 1 to 5.

9.37 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 9-35 shows the format of the *LLAddr* register; Table 9.43 describes the *LLAddr* register fields.

Figure 9-35 LLAddr Register Format

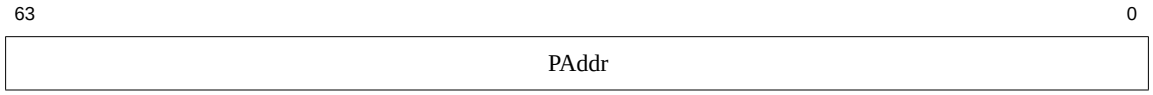

Table 9.43 LLAddr Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
PAddr	63..0	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

9.38 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 and DMTC0/DMFC0 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 and subsequent releases only) instruction whose address matches the Watch register address match conditions.

Figure 9-36 shows the format of the *WatchLo* register; Table 9.44 describes the *WatchLo* register fields.

Figure 9-36 WatchLo Register Format

Table 9.44 WatchLo Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
VAddr	63..3	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
I	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

Table 9.44 WatchLo Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
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
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

9.39 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 and DMTC0/DMFC0 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 9-37 shows the format of the *WatchHi* register; Table 9.45 describes the *WatchHi* register fields.

Figure 9-37 WatchHi Register Format

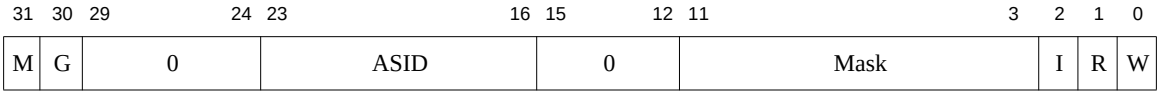

Table 9.45 WatchHi Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
M	31	If this bit is one, another pair of *WatchHi*/*WatchLo* registers is implemented at a MTC0 or MFC0 select field value of '*n+1*'	R	Preset	Required

Table 9.45 WatchHi Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
G	30	If this bit is one, any address that matches that specified in the *WatchLo* register will cause a watch exception. If this bit is zero, the *ASID* field of the *WatchHi* register must match the *ASID* field of the *EntryHi* register to cause a watch exception.	R/W	Undefined	Required
ASID	23..16	*ASID* value which is required to match that in the *EntryHi* register if the *G* bit is zero in the *WatchHi* register.	R/W	Undefined	Required
Mask	11..3	Optional bit mask that qualifies the address in the *WatchLo* 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. If this field is not implemented, writes to it must be ignored, and reads must return zero. Software may determine how many mask bits are implemented by writing ones the this field and then reading back the result.	R/W	Undefined	Optional
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	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)
W	0	This 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	29..24, 15..12	Must be written as zero; returns zero on read.	0	0	Reserved

9.40 XContext Register (CP0 Register 20, Select 0)

Compliance Level: *Required* for 64-bit TLB-based MMUs. *Optional* otherwise.

The *XContext* 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 *XContext* register is primarily intended for use with the XTLB Refill handler, but is also loaded by hardware on a TLB Refill. However, it is unlikely to be useful to software in the TLB Refill Handler. The *XContext* register duplicates some of the information provided in the *BadVAddr* register.

If $Config3_{CTXTC}$ =0 then the *XContext* register is organized in such a way that the operating system can directly reference a 16-byte structure in memory that describes the mapping. For PTE structures of other sizes, the content of this register can be used by the TLB refill handler after appropriate shifting and masking.

If $Config3_{CTXTC}$ =0 then a TLB exception (TLB Refill, XTLB Refill, TLB Invalid, or TLB Modified) causes bits 63..62 of the virtual address to be written into the R field and bits *SEGBITS*-1..13 of the virtual address to be written into the *BadVPN2* field of the *XContext* register. The *PTEBase* field is written and used by the operating system.

The *BadVPN2* and *R* fields of the *XContext* register are not defined after an address error exception and these fields may be modified by hardware during the address error exception sequence.

Figure 9-38 shows the format of the *XContext* register when $Config3_{CTXTC}$ =0; Table 9.46 describes the *XContext* register fields when $Config3_{CTXTC}$ =0. In Figure 9-38, bit numbers above the figure use the symbol *SEGBITS*; bit number under the figure assume that *SEGBITS* has the value 40.

Figure 9-38 XContext Register Format when $Config3_{CTXTC}$=0

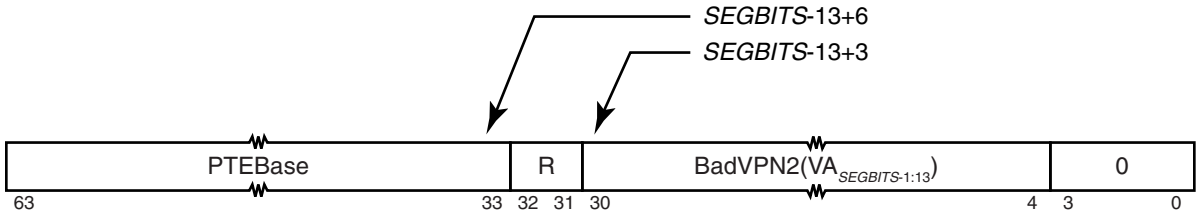

Table 9.46 XContext Register Fields when $Config3_{CTXTC}$=0

Field		Description	Read / Write	Reset State	Compliance
Name	Bits				
PTEBase	63 .. *SEGBITS*-13+6 (63..33 assuming *SEGBITS* is 40)	This field is for use by the operating system and is normally written with a value that allows the operating system to use the *XContext* Register as a pointer into the current PTE array in memory	R/W	Undefined	Required

Table 9.46 XContext Register Fields when $Config3_{CTXTC}$=0

Field		Description	Read / Write	Reset State	Compliance
Name	Bits				
R	*SEGBITS*-13+5 .. *SEGBITS*-13+4 (32..31 assuming *SEGBITS* is 40)	The *Region* field contains bits 63..62 of the virtual address. Encoding / Meaning: 0b00 – xuseg 0b01 – xsseg: supervisor address region. If Supervisor Mode is not implemented, this encoding is reserved 0b10 – Reserved 0b11 – xkseg For processors implementing $Config_{AT}$ = 1 (access to 32-bit compatibility segments only), only the 0b00 and 0b11 values are supplied by the processor on an exception.	R	Undefined	Required
BadVPN2	*SEGBITS*-13+3 .. 4 (30..4 assuming *SEGBITS* is 40)	The *Bad Virtual Page Number/2* field is written by hardware on a miss. It contains bits $VA_{SEGBITS-1..13}$ of the virtual address that missed.	R	Undefined	Required
0	3..0	Must be written as zero; returns zero on read.	0	0	Reserved

If $Config3_{CTXTC}$ =1 then the pointer implemented by the *XContext* register can point to any power-of-two-sized PTE structure within memory. This allows the TLB refill handler to use the pointer without additional shifting and masking steps. Depending on the value in the *XContextConfig* register, it may point to an 8-byte pair of 32-bit PTEs within a single-level page table scheme, or to a first level page directory entry in a two-level lookup scheme.

If $Config3_{CTXTC}$ =1 then the a TLB exception (Refill, Invalid, or Modified) causes bits $VA_{SEGBITS-1:SEGBITS-(X-Y)}$ to be written to a variable range of bits "(X-1):Y" of the *XContext* register, where this range corresponds to the contiguous range of set bits in the *XContextConfig* register. The exception causes bits 63..62 of the virtual address to be written into the R field. Bits 63:X+2 are R/W to software, and are unaffected by the exception. Bits Y-1:0 are unaffected by the exception. If X = 31 and Y = 4, i.e. bits 30:4 are set in *XContextConfig*, the behavior is identical to the standard MIPS III *XContext* register (bits 30:4 are filled with $VA_{39:13}$ when SEGBITS equals 40). Although the fields have been made variable in size and interpretation, the MIPS64 nomenclature is retained. Bits 63:X are referred to as the *PTEBase* and R fields, and bits X-1:Y are referred to as *BadVPN2*.

The value of the *XContext* register is **UNPREDICTABLE** following a modification of the contents of the *XContextConfig* register.

Figure 9-39 shows the format of the *XContext* Register when $Config3_{CTXTC}$ =1; Table 9.47 describes the *XContext* register fields $Config3_{CTXTC}$ =1.

Figure 9-39 XContext Register Format when Config3$_{CTXTC}$=1

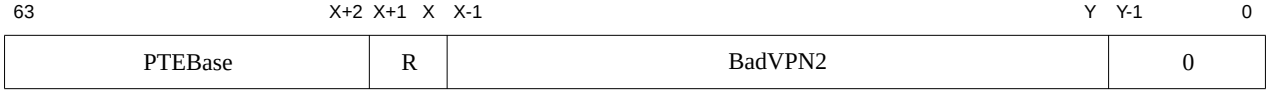

Table 9.47 XContext Register Field Descriptions when Config3$_{CTXTC}$=1

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
PTEBase	Variable, 63:X+2 where X in {63..0}. May be null.	This field is for use by the operating system and is normally written with a value that allows the operating system to use the *Context* Register as a pointer to an array of data structures in memory corresponding to the address region containing the virtual address which caused the exception.	R/W	Undefined	Required
R	X+1:X where X in {63..0}. May be null.	The *Region* field contains bits 63..62 of the virtual address. **Encoding** / **Meaning** 0b00 — xuseg 0b01 — xsseg: supervisor address region. If Supervisor Mode is not implemented, this encoding is reserved 0b10 — Reserved 0b11 — xkseg For processors implementing Config$_{AT}$ = 1 (access to 32-bit compatibility segments only), only the 0b00 and 0b11 values are supplied by the processor on an exception.	R	Undefined	Required
BadVPN2	Variable, (X-1):Y where X in {64..1} and Y in {63..0}. May be null.	This field is written by hardware on a TLB exception. It contains bits VA$_{SEGBITS-1:SEGBITS-(X-Y)}$ of the virtual address that caused the exception.	R	Undefined	Required
0	Variable, (Y-1):0 where Y in {63:1}. May be null.	Must be written as zero; returns zero on read.	0	0	Reserved

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

Compliance Level: *Implementation Dependent*.

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

9.42 Debug Register (CP0 Register 23, Select 0)

Compliance Level: *Optional*.

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

9.43 Debug2 Register (CP0 Register 23, Select 6)

Compliance Level: *Optional*.

The *Debug2* register is part of the EJTAG specification. Refer to that specification for the format and description of this register.

9.44 DEPC Register (CP0 Register 24)

Compliance Level: *Optional*.

The *DEPC* register is a read-write register that contains the address at which processing resumes after a debug exception has been serviced. It is part of the EJTAG specification and the reader is referred there for the format and description of the register. All bits of the *DEPC* register are significant and must be writable.

When a debug exception occurs, the processor writes the *DEPC* register with,

- 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.

The processor reads the *DEPC* register as the result of execution of the DERET instruction.

Software may write the *DEPC* register to change the processor resume address and read the *DEPC* register to determine at what address the processor will resume.

9.44.1 Special Handling of the DEPC Register in Processors That Implement the MIPS16e ASE or microMIPS64 Base Architecture

In processors that implement the MIPS16e ASE or the microMIPS64 base architecture, the *DEPC* register requires special handling.

When the processor writes the *DEPC* register, it combines the address at which processing resumes with the value of the *ISA Mode* register:

$$\text{DEPC} \leftarrow \text{resumePC}_{63..1} \parallel \text{ISAMode}_0$$

“resumePC” is the address at which processing resumes, as described above.

When the processor reads the *DEPC* register, it distributes the bits to the *PC* and *ISA Mode* registers:

$$\text{PC} \leftarrow \text{DEPC}_{63..1} \parallel 0$$
$$\text{ISAMode} \leftarrow \text{DEPC}_0$$

Software reads of the *DEPC* register simply return to a GPR the last value written with no interpretation. Software writes to the *DEPC* register store a new value which is interpreted by the processor as described above.

9.45 Performance Counter Register (CP0 Register 25)

Compliance Level: *Recommended*.

The 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 or 64-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 9.48 shows an example of two performance counters and how they map into the select values of the *PerfCnt* register.

Table 9.48 Example Performance Counter Usage of the PerfCnt CP0 Register

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

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 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 9-40 shows the format of the Performance Counter Control Register; Table 9.49 describes the Performance Counter Control Register fields.

Figure 9-40 Performance Counter Control Register Format

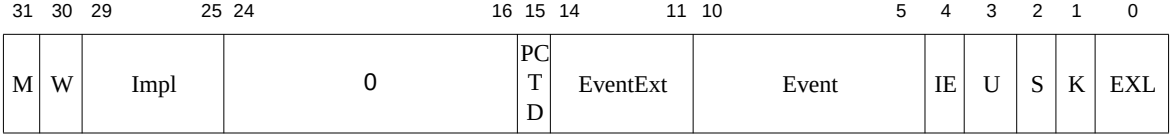

Table 9.49 Performance Counter Control Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>M</td><td>31</td><td>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’.</td><td>R</td><td>Preset</td><td>Required</td></tr>
<tr><td>W</td><td>30</td><td>Specifies the width of the corresponding Counter register, as follows:<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Width of the corresponding Counter register is 32 bits</td></tr><tr><td>1</td><td>Width of the corresponding Counter register is 64 bits</td></tr></table></td><td>R</td><td>Preset</td><td>Required (Release 2)</td></tr>
<tr><td>Impl</td><td>29:25</td><td>This field is implementation dependent and is not specified by the architecture.

If not used by the implementation, must be written as zero; returns zero on read.</td><td></td><td>Undefined

0 if not used by the implementation</td><td>Optional</td></tr>
<tr><td>0</td><td>24..16</td><td>Must be written as zero; returns zero on read</td><td>0</td><td>0</td><td>Reserved</td></tr>
<tr><td>PCTD</td><td>15</td><td>Performance Counter Trace Disable.
The PDTrace facility (revision 6.00 and higher) has the ability to trace Performance Counter in its output. This bit is used to disable the specified performance counter from being traced when performance counter trace is enabled and a performance counter trace event is triggered.<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Tracing is enabled for this counter.</td></tr><tr><td>1</td><td>Tracing is disabled for this counter.</td></tr></table></td><td>RW</td><td>0</td><td>Required if PDTrace Performance Counter Tracing feature is implemented.</td></tr>
<tr><td>EventExt</td><td>14..11</td><td>In some implementations which support more than the the 64 encodings possible in the 6-bit Event field, the EventExt field acts as an extension to the Event field. In such instances the event selection is the concatentation of the two fields, i.e., EventExt|Event.

The actual field width is implementation dependent. Any bits that are not implemented read as zero and are ignored on write.</td><td>RW</td><td>Undefined</td><td>Optional</td></tr>
</table>

Table 9.49 Performance Counter Control Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>Event</td><td>10..5</td><td>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</td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>IE</td><td>4</td><td>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 Status register.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Performance counter interrupt disabled</td></tr><tr><td>1</td><td>Performance counter interrupt enabled</td></tr></table></td><td>R/W</td><td>0</td><td>Required</td></tr>
<tr><td>U</td><td>3</td><td>Enables event counting in User Mode. Refer to Section 3.4 “User Mode” on page 20 for the conditions under which the processor is operating in User Mode.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Disable event counting in User Mode</td></tr><tr><td>1</td><td>Enable event counting in User Mode</td></tr></table></td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>S</td><td>2</td><td>Enables event counting in Supervisor Mode (for those processors that implement Supervisor Mode). Refer to Section 3.3 “Supervisor Mode” on page 20 for the conditions under which the processor is operating in Supervisor mode.
If the processor does not implement Supervisor Mode, this bit must be ignored on write and return zero on read.
<table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Disable event counting in Supervisor Mode</td></tr><tr><td>1</td><td>Enable event counting in Supervisor Mode</td></tr></table></td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
</table>

Table 9.49 Performance Counter Control Register Field Descriptions

<table>
<tr><th colspan="2">Fields</th><th rowspan="2">Description</th><th rowspan="2">Read / Write</th><th rowspan="2">Reset State</th><th rowspan="2">Compliance</th></tr>
<tr><th>Name</th><th>Bits</th></tr>
<tr><td>K</td><td>1</td><td>Enables event counting in Kernel Mode. Unlike the usual definition of Kernel Mode as described in Section 3.2 “Kernel Mode” on page 19, this bit enables event counting only when the EXL and ERL bits in the Status register are zero.
<table>
<tr><th>Encoding</th><th>Meaning</th></tr>
<tr><td>0</td><td>Disable event counting in Kernel Mode</td></tr>
<tr><td>1</td><td>Enable event counting in Kernel Mode</td></tr>
</table></td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
<tr><td>EXL</td><td>0</td><td>Enables event counting when the EXL bit in the Status register is one and the ERL bit in the Status register is zero.
<table>
<tr><th>Encoding</th><th>Meaning</th></tr>
<tr><td>0</td><td>Disable event counting while EXL = 1, ERL = 0</td></tr>
<tr><td>1</td><td>Enable event counting while EXL = 1, ERL = 0</td></tr>
</table>
Counting is never enabled when the ERL bit in the Status register or the DM bit in the Debug register is one.</td><td>R/W</td><td>Undefined</td><td>Required</td></tr>
</table>

The Counter Register associated with each performance counter increments once for each enabled event. Figure 9-41 shows the format of the Performance Counter Counter Register; Table 9.50 describes the Performance Counter Counter Register fields.

Figure 9-41 Performance Counter Counter Register Format

31 or 63 ... 0
Event Count

Table 9.50 Performance Counter Counter Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
Event Count	31..0 or 63..0	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 *Cause* register. The width of the counter is either 32 bits or 64 bits depending on the value of the *W* bit in the corresponding Performance Counter Control Register.	R/W	Undefined	Required

Programming Note:

In Release 2 of the Architecture, the EHB instruction can be used to make interrupt state changes visible when the IE field of the Control register or the Event Count Field of the Counter register are written. See sECTION 6.1.2.1 “Software Hazards and the Interrupt System” on page 58.

9.46 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.

9.47 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.

9.48 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.

9.49 DataLo Register (CP0 Register 28, Select 1, 3)

Compliance Level: *Optional*.

The *DataLo* and *DataHi* registers are 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.

9.50 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.

9.51 DataHi Register (CP0 Register 29, Select 1, 3)

Compliance Level: *Optional*.

The *DataLo* and *DataHi* registers are 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.

9.52 ErrorEPC (CP0 Register 30, Select 0)

Compliance Level: Required.

The *ErrorEPC* register is a read-write register, similar to the *EPC* register, at which processing resumes after a Reset, Soft Reset, Nonmaskable Interrupt (NMI) or Cache Error exceptions (collectively referred to as error exceptions). Unlike the *EPC* register, there is no corresponding branch delay slot indication for the *ErrorEPC* register. All bits of the *ErrorEPC* register are significant and must be writable.

When an error exception occurs, the processor writes the *ErrorEPC* register with:

- 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.

The processor reads the *ErrorEPC* register as the result of execution of the ERET instruction.

Software may write the *ErrorEPC* register to change the processor resume address and read the *ErrorEPC* register to determine at what address the processor will resume

Figure 9-42 shows the format of the *ErrorEPC* register; Table 9.51 describes the *ErrorEPC* register fields.

Figure 9-42 ErrorEPC Register Format

Table 9.51 ErrorEPC Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
ErrorEPC	63..0	Error Exception Program Counter	R/W	Undefined	Required

9.52.1 Special Handling of the ErrorEPC Register in Processors That Implement the MIPS16e ASE or microMIPS64 Base Architecture

In processors that implement the MIPS16e ASE or microMIPS64 base architecture, the *ErrorEPC* register requires special handling.

When the processor writes the *ErrorEPC* register, it combines the address at which processing resumes with the value of the *ISA Mode* register:

$$\texttt{ErrorEPC} \leftarrow \texttt{resumePC}_{63..1} \parallel \texttt{ISAMode}_{0}$$

“resumePC” is the address at which processing resumes, as described above.

When the processor reads the *ErrorEPC* register, it distributes the bits to the *PC* and *ISAMode* registers:

```
PC ← ErrorEPC63..1 || 0
ISAMode ← ErrorEPC0
```

Software reads of the *ErrorEPC* register simply return to a GPR the last value written with no interpretation. Software writes to the *ErrorEPC* register store a new value which is interpreted by the processor as described above.

9.53 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.

The *DESAVE* register is meant to be used solely while in Debug Mode. If kernel mode software uses this register, it would conflict with debugging kernel mode software. For that reason, it is strongly recommended that kernel mode software not use this register. If the *KScratch** registers are implemented, kernel software can use those registers.

9.54 KScratch*n* Registers (CP0 Register 31, Selects 2 to 7)

Compliance Level: *Optional, KScratch1 and KScratch2 at selects 2, 3 are recommended.*

The *KScratchn* registers are read/write registers available for scratch pad storage by kernel mode software. These registers are 32bits in width for 32-bit processors and 64bits for 64-bit processors.

The existence of these registers is indicated by the KScrExist field within the *Config4* register. The KScrExist field specifies which of the selects are populated with a kernel scratch register.

Debug Mode software should not use these registers, instead debug software should use the DESAVE register. If EJTAG is implemented, select 0 should not be used for a KScratch register. Select 1 is being reserved for future debug use and should not be used for a KScratch register.

Figure 9-43 KScratch*n* Register Format

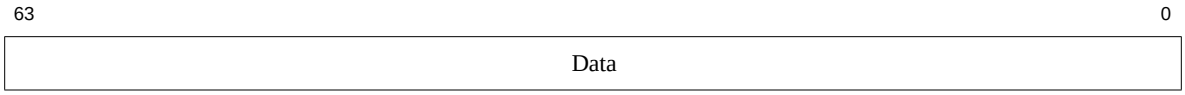

Table 9.52 KScratch*n* Register Field Descriptions

Fields		Description	Read / Write	Reset State	Compliance
Name	Bits				
Data	63:0	Scratch pad data saved by kernel software.	R/W	Undefined	Optional

Appendix A

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 MIPS64/microMIPS64 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. It is not anticipated that MIPS64 processors that implement a fixed-mapping MMU will require a 64-bit address capability. As a result, the description below is given assuming a 32-bit address.

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 A.1 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.

Table A.1 Physical Address Generation from Virtual Addresses

Segment Name	Virtual Address	Generates Physical Address	
		$Status_{ERL}$ = 0	$Status_{ERL}$ = 1
useg suseg kuseg	0x0000 0000 through 0x7FFF FFFF	0x4000 0000 through 0xBFFF FFFF	0x0000 0000 through 0x7FFF FFFF
kseg0	0x8000 0000 through 0x9FFF FFFF	0x0000 0000 through 0x1FFF FFFF	

Table A.1 Physical Address Generation from Virtual Addresses (Continued)

Segment Name	Virtual Address	Generates Physical Address	
		$Status_{ERL}$ = 0	$Status_{ERL}$ = 1
kseg1	0xA000 0000 through 0xBFFF FFFF	0x0000 0000 through 0x0x1FFF FFFF	
sseg ksseg kseg2	0xC000 0000 through 0xDFFF FFFF	0xC000 0000 through 0xDFFF FFFF	
kseg3	0xE000 0000 through 0xFFFF FFFF	0xE000 0000 through 0xFFFF FFFF	

Note that this mapping means that physical addresses 0x2000 0000 through 0x3FFF FFFF are inaccessible when the ERL bit is off in the *Status* register, and physical addresses 0x8000 0000 through 0xBFFF FFFF are inaccessible when the ERL bit is on in the *Status* register.

Figure A-1 shows the memory mapping when the ERL bit in the *Status* register is zero; Figure A-2 shows the memory mapping when the ERL bit is one.

Figure A-1 Memory Mapping when ERL = 0

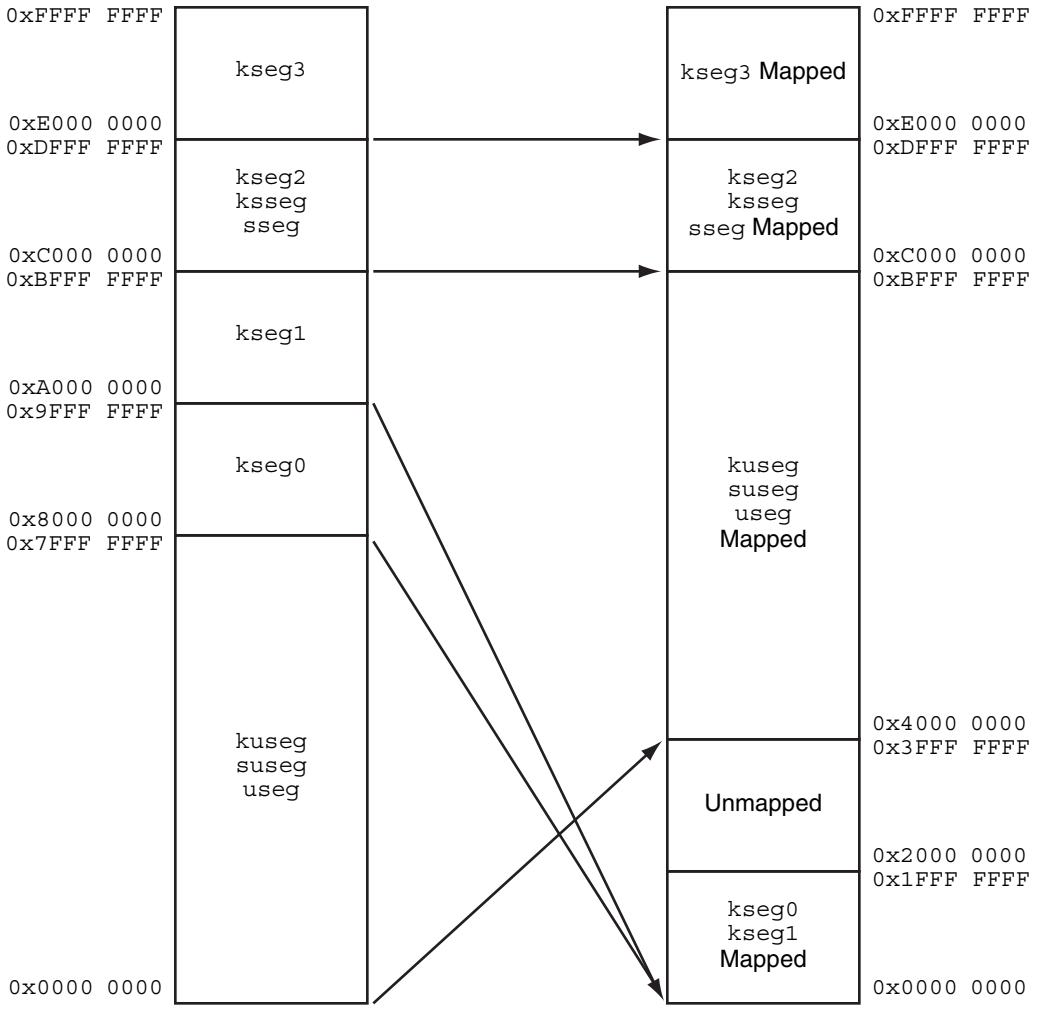

Figure A-2 Memory Mapping when ERL = 1

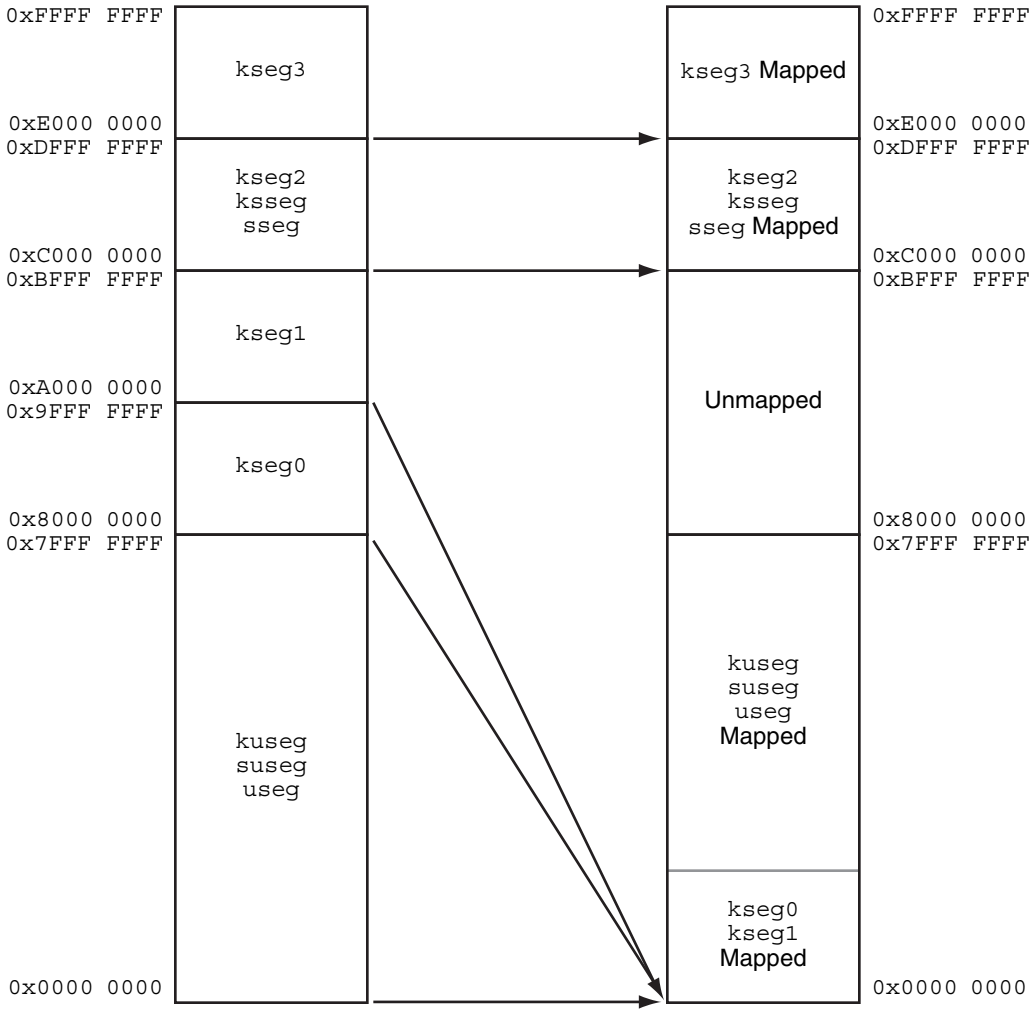

A.1.2 Cacheability Attributes

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 A-3 shows the format of the additions to the *Config* register; Table A.2 describes the new *Config* register fields.

Figure A-3 Config Register Additions

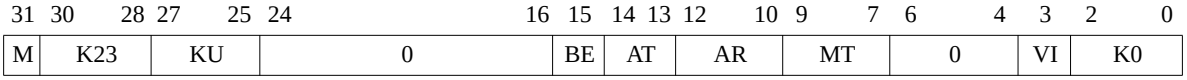

Table A.2 Config Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
K23	30:28	Kseg2/Kseg3 cacheability and coherency attribute. See Table 9.9 on page 106 for the encoding of this field.	R/W	Undefined	Required
KU	27:25	Kuseg cacheability and coherency attribute when $Status_{ERL}$ is zero. See Table 9.9 on page 106 for the encoding of this field.	R/W	Undefined	Required

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 effects of a read or write to these registers are **UNDEFINED**.
- The TLBWR, TLBWI, TLBP, and TLBR instructions are no longer required and must 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 A-4 shows the logical arrangement of a BAT entry.

Figure A-4 Contents of a BAT Entry

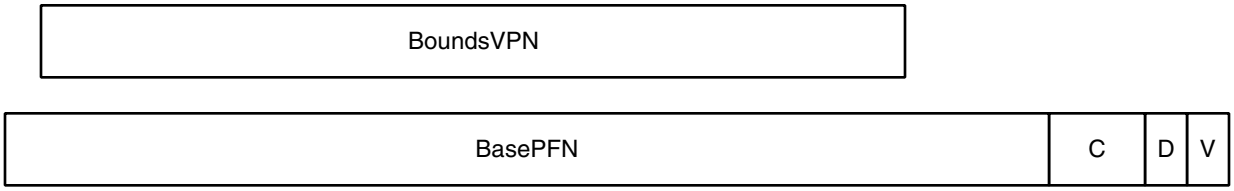

The BAT is indexed by the reference type and the address region to be checked as shown in Table A.3.

Table A.3 BAT Entry Assignments

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

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:

```
i ← SelectIndex (reftype, va)
bounds ← BAT[i]BoundsVPN || 1^12
pfn ← BAT[i]BasePFN
c ← BAT[i]C
d ← BAT[i]D
v ← BAT[i]V
if (va > bounds) or (v = 0) then
    InitiateTLBInvalidException(reftype)
endif
if (d = 0) and (reftype = store) then
    InitiateTLBModifiedException()
```

```
endif
pa ← 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 signal a Reserved Instruction Exception.

A.3 Dual Variable-Page-Size and Fixed-Page-Size TLBs

Most MIPS CPU cores implement a fully associative Joint TLB. Unfortunately, such fully-associative structures can be slow, can require a large amount of logic components to implement and can dissipate a lot of power. The number of entries for a fully associative array that can be practically implemented is not large.

In high performance systems, it is desirable to minimize the frequency of TLB misses. In small and low-cost systems, it is desirable to keep the implementation costs of a TLB to a minimum. This section describes an optional alternative MMU configuration which decreases the implementation costs of a small TLB as well as allows for a TLB that can map a very large number of pages to be reasonably implemented.

A.3.1 MMU Organization

This alternative MMU configuration uses two TLB structures.

1. This first TLB is called the Fixed-Page-Size TLB or the FTLB.
    - At any one time, all entries within the FTLB use a shared, common page size.
    - The FTLB is not fully-associative, but rather set associative.
    - The number of ways per set is implementation specific.
    - The number of sets is implementation specific.
    - The common page size is also implementation specific.
    - The common page size is allowed to be software configurable. The choice of the common page size is done once for the entire FTLB, not on a per-entry basis. This configuration by software can only be done after a full flush/initialization of the FTLB, before there are any valid entries within the FTLB. Implementations are also allowed to support only one page size for the FTLB - in that case, the FTLB page size is fixed by hardware and not software configurable.
2. The second TLB is called the Variable-Page-Size TLB or the VTLB.
    - The choice of page size is done on a per-entry basis. That is, one VTLB entry can use a pagesize that is different from the size used by another VTLB entry.
    - The VTLB is fully-associative.
    - The number of entries is implementation specific.
    - The set of allowable page sizes for VTLB entries is implementation specific.

Just as for the JTLB, both the FTLB and VTLB are shared between the instruction stream and the data stream. For address translation, the virtual address is presented to both the FTLB and VTLB in parallel. Entries in both structures are accessed in parallel to search for the physical address.

The use of two TLB structures has these benefits:

- The implementation costs of building a set-associative TLB with many entries can be much less than that of implementing a large fully-associative TLB.

- The existence of a VTLB retains the capability of using large pages to map large sections of physical memory without consuming a large number of entries in the FTLB.

Random replacement of pages in the MMU happens mainly in the FTLB. In most operating systems, on-demand paging only uses one page size so the FTLB is sufficient for this purpose. Some of the address bits of the specified virtual address are used to index into the FTLB as appropriate for the chosen FTLB array size. The method of choosing which FTLB way to modify is implementation specific.

The VTLB is very similiar to the JTLB. The *C0_PageMask* register is used to program the page size used for a particular VTLB entry.

The configuration of the FTLB is reflected in the FTLB fields within the new *C0_Config4* register. The size of the VTLB is reflected in the $C0\_Config1_{MMUSize-1}$ field. The presense of the dual FTLB and VTLB is denoted by the value of 0x4 in $C0\_Config_{MT}$ register field. These registers are described in "Changes to the COP0 Registers" on page 226.

Most implementations would choose to build a VTLB with a smaller number of entries and a FTLB with a larger number of entries. This combination allows for many on-demand fixed-sized pages as well as for a small number of large address blocks to be simultaneously mapped by the MMU.

A.3.2 Programming Interface

The software programming interface used for the fully-associative JTLB is maintained as much as possible to decrease the amount of software porting.

Also for that purpose, each entry in the FTLB as well as the VTLB use one tag (VPN2) to map two physical pages (PFN), just as in the JTLB. The entries in either array are accessed through the *C0_EntryHi* and *C0_EntryLo0/1* registers.

Entries in either array (FTLB or VTLB) can be accessed with the TLBWI and TLBWR instructions.

The PageMask register is used to set the page size for the VTLB entries. This register is also used to choose which array (FTLB or VTLB) to write for the TLBWR instruction.

For the rest of this section, the following parameters are used:

3. FPageSize - the page size used by the FTLB entries

4. FSetSize - Number of entries in one way of the FTLB.

5. FWays - Number of ways within a set of the FTLB.

6. VIndex - Number of entries in the VTLB.

For the *C0_Index* , the *C0_Wired* registers, the TLBP, TLBR and TLBWI instructions; the VTLB occupies indices 0 to VIndex-1. The FTLB occupies indices VIndex to VIndex + (FSetSize * FWays)-1.

The TLBP instruction produces a value which can be used by the TLBWI instruction without modification by software. When referring to the FTLB, the value is the concatentation of the selected FTLB way and set, and incremented by the size of the VTLB. For example, {selected FTLB Way, selected FTLB Set} + VIndex.

If *C0_PageMask* is set to the page size used by the FTLB, the TLBWR instruction modifies entries within the FTLB.

How the FTLB set-associative array is indexed is implementation specific. In any indexing scheme, the least significant address bit that can be used for indexing is $\log_2$(FPageSize)+1. The number of index bits needed to select the correct set within the FTLB array is $\log_2$(FSetSize).

Since the FTLB array can be modified through the TLBWI instruction, it is possible for software to choose an inappropriate FTLB index value for the specified virtual address. In this case, it is implementation specific whether a Machine Check exception is generated for the TLBWI instruction.

The method of choosing which FTLB way to modify is implementation specific.

If *C0_PageMask* is not set to the pagesize used by the FTLB, the TLBWR instruction modifies entries within the VTLB. The VTLB entry to be written is specified by the $\log_2$(VIndex) least significant bits of the *C0_Random* register value.

For both the TLBWR and TLBWI instruction, it is implementation specific whether both (FTLB and VTLB) arrays are checked for duplicate or overlapping entries and whether a Machine Check exception is generated for these cases.

A.3.2.1 Example with chosen FTLB and VTLB sizes

As an example, let's assume an implementation chooses these values:

1. FPageSize - 4KB used by the FTLB entries
2. FSetSize - 128 in one way of the FTLB.
3. FWays - 4 ways within a set of the FTLB. (The FTLB has (128 sets x 4 ways/set) 512 entries, capable of mapping (512 entries x 2 pages/entry x 4KB/page) 4MB of address space simultaneously.
4. VIndex - 8 entries in the VTLB.

For the *C0_Index*, the *C0_Wired* registers, the TLBP, TLBR and TLBWI instructions; the VTLB occupies indices 0 to 7. The FTLB occupies indices 8 to 519.

The FTLB entries have a VPN2 field which starts at virtual address bit 12.

The least significant virtual address bit that can be used for FTLB indexing is virtual address 13. To index the FTLB set-associative array, 7 index bits are needed.

In this simple example, the design uses contiguous virtual address bits directly for indexing the FTLB (it does not create a hash for the FTLB indexing). The FTLB set-associative array is indexed using virtual address bits 19:13. The TLBWR instruction uses these address bits held in *C0_EntryHi.*

In this simple example, the design uses a cycle counter of 2 bits for way selection within the FTLB.

The Random register field within *C0_Random* is 3 bits wide to select the entry within the VTLB.

A.3.3 Changes to the TLB Instructions

TLBP

Both the VTLB and the FTLB are probed in parallel for the specified virtual address.

If the address hits in the VTLB, *C0_Index* specifies the entry within the VTLB [a value within 0 to VIndex-1].

If the address hits in the FTLB, *C0_Index* specifies the entry within the FTLB [a value within VIndex to VIndex+(FSetSize * FWays)-1]. Which bits are used to encode the selected FTLB set as opposed to which bits are used to encode the selected FTLB way is implementation specific, but must match what is expected by the TLBWI instruction implementation. *C0_PageMask* reflects the pagesize used by the FTLB.

TLBR

Either a VTLB entry or a FTLB entry is read depending on the specified index in *C0_Index*.

Index values of 0 to VIndex-1 access the VTLB. Index values VIndex to VIndex+(FSetSize * FWays)-1 access the FTLB.

TLBWI

Either the VTLB or FTLB entry is written depending on the specified index in *C0_Index*.

Index values of 0 to VIndex-1 access the VTLB. Index values VIndex to VIndex+(FSetSize * FWays)-1 access the FTLB.

It is implementation specific if the hardware checks the VPN2 field of *C0_EntryHi* is appropriate for the specified set within the FTLB. The implementation may generate a machine-check exception if the VPN2 field is not appropriate for the specified set.

It is implementation specific if the hardware checks both arrays (FTLB and VTLB) for valid duplicate or overlapping entries and if the hardware signals a Machine Check exception for these cases.

TLBWR

Either the VTLB or FTLB entry is written depending on the specified pagesize in C0_PageMask.

If *C0_PageMask* is set to any pagesize other than that used by the FTLB, the TLBWR instruction modifies a VTLB entry. The VTLB entry is specified by the Random register field within *C0_Random*.

If *C0_PageMask* is set to the pagesize used by the FTLB, the TLBWR modifies a FTLB entry. The FTLB set-associative array is indexed in an implementation-specific manner.

The method of selecting which FTLB way to modify is implementation specific.

It is implementation specific if the hardware checks both arrays (FTLB and VTLB) for valid duplicate or overlapping entries and if the hardware signals a Machine Check exception for these cases.

A.3.4 Changes to the COP0 Registers

C0_Config4 (CP0 Register 16, Select 4)

A new register introduced to reflect the FTLB configuration. *Config4*$_{MMUExtDef}$ register field must be set to a value of 2 to reflect that the Dual VTLB and FTLB configuration is implemented. If either *Config4* is not implemented or the *Config4*$_{MMUExtDef}$ field is not fixed to 2, the Dual VTLB/FTLB configuration is not implemented.

If *Config4*$_{MMUExtDef}$ is fixed to a value of 2, the FTLBPageSize, FTLBWays and FTLBSets fields reflect the FTLB configuration. Please refer to "Configuration Register 4 (CP0 Register 16, Select 4)" on page 183 for more detail on this register.

C0_Config1 (CP0 Register 16, Select 1)

If *Config4*$_{MMUExtDef}$ is fixed to a value of 2, the MMUSize-1 register field is redefined to reflect only the size of the VTLB.

C0_Config (CP0 Register 16, Select 0)

If *Config4*$_{MT}$ is fixed to a value of 4, the implemented MMU Type is the dual FTLB and VTLB configuration.

C0_Index (CP0 Register 0, Select 0)

If *Config4*$_{MMUExtDef}$ is fixed to a value of 2, the register is redefined in this way:

The value held in the Index field can refer to either an entry in the FTLB or the VTLB. Index values of 0 to VIndex-1 access the VTLB. Index values VIndex to VIndex+(FSetSize * FWays)-1 access the FTLB. Which bits in the register field which encode the FTLB set as opposed to which bits encode the FTLB way is implementation specific, but must match what is expected by the TLBWI instruction implementation.

C0_Random (CP0 Register 1, Select 0)

If *Config4*$_{MMUExtDef}$ is fixed to a value of 2, the register is redefined in this way:

If the value in *C0_PageMask* is not set to the page-size used by the FTLB, and a TLBWR instruction is executed, a VTLB entry is modified. The Random register field is used to select the VTLB entry which is modified.

If the value in *C0_PageMask* is set to the page-size used by the FTLB, and a TLBWR instruction is executed, a FTLB entry is modified. It is implementation specific whether the *C0_RANDOM* register is used to select the FTLB entry.

The upper bound of the Random register field value is VIndex.

C0_Wired (CP0 Register 6, Select 0)

If *Config4*$_{MMUExtDef}$ is fixed to a value of 2, the Wired register field can only hold a value of VIndex-1 or less. That is, only VTLB entries can be wired down.

C0_PageMask (CP0 Register 5, Select 0)

If $Config4_{MMUExtDef}$ is fixed to a value of 2, the register is redefined in this way:

The Mask and MaskX field values determine whether the VTLB or the FTLB is modified by a TLBWR instruction.

The Mask and MaskX register fields do not affect the TLB address match operation for FTLB entries. The pagesize used by the FTLB entries are specified by the $Config4_{FPageSize}$ register field.

The software writeable bits in the Mask and MaskX fields reflect what page sizes are available in the VTLB. These fields do not reflect the page sizes which are available in the FTLB.

A.3.5 Software Compatibility

One of the main software visible changes introducted by this alternative MMU are the values reported in the *C0_Index* register. Previously, it was just a simple linear index. For this alternative MMU configuration, the value reflects both a selected way as well as a selected set when a FTLB entry is specified.

Fortunately, this Index value isn't frequently generated by software nor read by software. Instead, the contents of the *C0_Index* register is generated by hardware upon a TLBP instruction. Software then just issues the TLBWI instruction once the *C0_EnLo** registers have been appropriately modified.

Another software visible change is that the MMUSize-1 field no longer reports the entire MMU size. For TLB initialization and TLB flushing, the contents of $Config1_{MMUSize-1}$, $Config4_{FTLBWays}$ and $Config4_{FTLBSets}$ register fields must all be read to calculate the entire number of TLB entries that must be initialized. TLB initialization and flushing are the only times software needs to generate an Index value to write into the *C0_Index* register.

Only the VTLB entries may be wired down. This limitation is due to using some of the *EntryHi* VPN2 bits to index the FTLB array.

If a page using the FTLB page-size is to be wired down, that page must be programmed into the VTLB using the TLBWI instruction, as the TLBWR instruction would only access the FTLB in that situation and could not access any wired-down TLB entry. The TLBWI instruction is normally used for wired-down pages, so this restriction should not affect existing software.

Appendix B

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 documentation.
0.95	March 12, 2001	Clean up document for external review release
1.00	August 29, 2002	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 $Config_{VI}$ to denote an instruction cache with both virtual indexing and virtual tags. • 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 4..3 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. • Modify the description of PageMask and the TLB lookup process to acknowledge the fact that not all implementations may support all page sizes.
1.90	September 1, 2002	Update the specification with the changes introduced in Release 2 of the Architecture. Changes in this revision include: • The following new Coprocessor 0 registers were added: EBase, HWREna, IntCtl, PageGrain, SRSCtl, SRSMap. • The following Coprocessor 0 registers were modified: Cause, Config, Config2, Config3, EntryHi, EntryLo0, EntryLo1, PageMask, PerfCnt, Status, WatchHi, WatchLo. • 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.

Revision	Date	Description
2.00	June 9, 2003	Complete the update to include Release 2 changes. These include: • Make bits 12..11 of the PageMask register power up zero and be gated by 1K page enable. This eliminates the problem of having these bits set to 0b11 on a Release 2 chip in which kernel software has not enabled 1K page support. • Correct the address of the cache error vector when the BEV bit is 1. It should be 0xBFC0.0300,. not 0xBFC0.0200. • Correct the introduction to shadow registers to note that the SRSCtl register is not updated at the end of an exception in which $Status_{BEV} = 1$. • 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 $EBase_{15..12}$ 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.
2.50	July 1, 2005	Changes in this revision: • Correct the encoding table description for the $Cause_{PCI}$ bit to indicate that the bit controlls the performance counter, not the timer interrupt. • Correct the figure Interrupt Generation for External Interrupt Controller Interrupt Mode to show $Cause_{IP1..0}$ going to the EIC, rather than $Status_{IP1..0}$ • Update all files to FrameMaker 7.1. • Update reset exception list to reflect missing Release 2 reset requirements. • Define bits 31..30 in the *HWREna* register as access enables for the implementation-dependent hardware registers 31 and 30. • Add definition for Coprocessor 0 Enable to Operating Modes chapter. • Add K23 and KU fields to main Config register definition as a pointer to the Fixed Mapping MMU appendix. • Add specific note about the need to implement all shadow sets between 0 and HSS - no holes are allowed. • Change the hazard from a software write to the $SRSCtl_{PSS}$ field and a RDPGPR and WRPGPR and instruction hazard vs. an execution hazard. • Correct the pseudo-code in the cache error exception description to reflect the Release 2 change that introduced EBase. • Document that EHB clears instruction state change hazards for writes to interrupt-related fields in the *Status*, *Cause*, *Compare*, and *PerfCnt* registers. • Note that implementation-dependent bits in the *Status* and *Config* registers should be defined in such a way that standard boot software will run, and that software which preserves the value of the field when writing the registers will also run correctly. • With Release 2 of the Architecture the FR bit in the *Status* register should be a R/W bit, not a R bit. • Improve the organization of the CP0 hazards table, and document that DERET, ERET, and exceptions and interrupts clear all hazards before the instruction fetch at the target instruction. • Add list of MIPS® MT CP0 registers and MIPS MT and MIPS® DSP present bits in the *Config3* register.

Revision	Date	Description
2.60	Jun 25, 2008	Changes in this revision: • Add the *UserLocal* register and access to it via the RDHWR instruction. • Operating Modes - footnote about ksseg/sseg • COP3 no longer usable for customer extensions • EIC Mode allows VectorNum != RIPL • CP0Regs Table - added missing EJTAG & PDTrace Registers • *C0_DataLo/Hi* are actually R/W • Hazards table - added a bunch of missing ones • Various typos fixed.
2.61	August 01, 2008	• In the *Status* register description, the ERL behavior description was incorrect in saying only 29bits of kuseg becomes uncached&unmapped.
2.62	January 2,009	• CCRes is accessed through $3 not $4 - *HWENA* register affected. • PCTD bit added to *C0_PerfCtl*.
2.70	January 22, 2009	MIPS Technologies-only release for internal review: • Added BigPages feature - Pages larger than 256MB are supported. *C0_PageMask* and *C0_Config3* affected. • Added CP0 Reg 31, Select 2 & 3 as kernel scratch registers. • Added VTLB/FTLB optional MMU configuration to Appendix A and *Config4* register for these new MMU configurations • Added CDMM chapter, *CDMMBase* COP0 Register, CDMM bit in *C0_Config3*, FDCI bit in *C0_Cause* register and IPFDC field in *IntCtl* register.
2.71	January 28, 2009	MIPS Technologies-only release for internal review: • EIC mode - revision 2.70, was actually missing the new option of EIC driving an explicit vector offset (not using VectorNumbers). • Clarified the text and diagrams for the 3 EIC options - RIPL=VectorNum, Explicit VectorNum; Explicit VectorOffset.
2.72	April 20, 2009	MIPS Technologies-only release for internal review: • Table was incorrectly saying ECR_{ProbEn} selected debug exception Vector. Changed to $ECR_{ProbTrap}$. • Added MIPS Technologies traditional meanings for CCA values. • Added list of COP2 instruction to COPUnusable Exception description. • Added statement that only uncached access is allowed to CDMM region. • Updated Exception Handling Operation pseudo-code for EIC Option_3 (EIC sends entire vector).
2.73	April 22, 2009	MIPS Technologies-only release for internal review: • Fixed comments for ASE.
2.74	June 03, 2009	MIPS Technologies-only release for internal review: • Added CDMM Enable Bit in *CDMMBase* COP0 register • Reserved CCA values can be used to init TLB; just can't be used for mapping.
2.75	June 12, 2009	MIPS Technologies-only release for internal review: • CDMMBase_Upper_Address Field doesn't have a fixed reset value. • Added DSP State Disabled Exception to *C0_Cause* Exception Type table.
2.80	July 20, 2009	• FTLB and VTLB MMU configuration denoted by 0x4 in $Config_{MT}$ • Added TLBP -> TLBWI hazard • Added KScrExist field in *Config4*.

Revision	Date	Description
2.81	September 22, 2009	MIPS Technologies-only release for internal review: • ContextConfig Register description added. • Context Register description updated for SmartMIPS behavior. • EntryLo* register descriptions updated for RI & XI bits. • TLB description and pseudo-code updated for RI & XI bits. • PageMask register updated for RIE and XIE bits. • Config3 register updated for CTXTC and RXI bits. • Reserve MCU ASE bits in C0_Cause and C0_Status. • Clean up description for KScratch registers - selects 2&3 are recommended, but additional scratch registers are allowed.
2.82	January 19, 2010	MIPS Technologies-only release for internal review: • Added Debug2 register.
3.00	March 8, 2010	• RI/XI feature moved from SmartMIPS ASE. • microMIPS features added • MCU ASE features added. • XI and RI exceptions can be programmed to use their own exception codes instead of using TLBL code. • XI and RI can be independently implemented as XIE and RIE bits are allowed to be Read-Only. • TCOpt Register added to C0 Register list. • Added encoding (0x7) for 32 sets for one cache way.
3.05	July 07, 2010	• CMGCRBase register added. • Lower bits of C0_Context register allowed to be write-able if Config3.CTXTC=1 and Config3.SM=0.
3.10	July 27, 2010	• Add XContextConfig register. • Explain the limits of the BadVPN2 field within Context and XContext registers and the relationships with the writeable bits within ContextConfig and XContextConfig registers.
3.11	April 24, 2011	MIPS Technologies-only release for internal review: • FPR registers are UNPREDICTABLE after change of Status.FR bit. • 1004K did not support CCA=0 • Config4 - KScratch Registers, mention that select 1 is reserved for future debugger use. • Context Register - the bit subscripts describing which VA bits go into the BadVPN2 field was incorrect for the case when the ContextConfig register is used. The correct VA bits are 31:31-((X-Y)-1) for MIPS32, 63:63-((X-Y)-1) for MIPS64.
3.12	April 28, 2011	• XContext & XContextConfig registers - be more explicit of the SEGBITS limitations. • ContextConfig Register is only 32-bits in width to be more compatible to MIPS32.