



# Intel® Itanium™ Processor Reference Manual for Software Development

---

Revision 2.0

*December 2001*





THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel® processors based on the Itanium architecture may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at <http://www.intel.com>.

Intel, Itanium, Pentium, VTune and MMX are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright © 2000-2001, Intel Corporation

\*Other names and brands may be claimed as the property of others.



# Contents

---

<b>1</b>	<b>About this Manual .....</b>	<b>1</b>
1.1	Overview of the Intel® Itanium™ Processor Reference Manual for Software Development .....	1
1.2	Terminology.....	1
1.3	Related Documents .....	2
1.4	Revision History .....	3
<b>2</b>	<b>Register Stack Engine Support.....</b>	<b>5</b>
2.1	RSE Modes .....	5
2.2	RSE and Clean Register Stack Partitions .....	5
<b>3</b>	<b>Virtual Memory Management Support.....</b>	<b>7</b>
3.1	Page Size Supported .....	7
3.2	Physical and Virtual Addresses .....	7
3.3	Region Register ID .....	7
3.4	Protection Key Register.....	7
<b>4</b>	<b>Processor Specific Write Coalescing (WC) Behavior .....</b>	<b>9</b>
4.1	Write Coalescing .....	9
4.2	WC Buffer Eviction Conditions .....	9
4.3	WC Buffer Flushing Behavior .....	10
<b>5</b>	<b>Model Specific Instruction Implementation .....</b>	<b>11</b>
5.1	ld.bias .....	11
5.2	lfetch Exclusive Hint .....	11
5.3	fwb.....	11
5.4	thash.....	12
5.5	ttag .....	12
5.6	ptc.e.....	13
5.7	mf.a .....	13
5.8	Prefetch Behavior .....	13
5.9	Temporal and Non-temporal Hints Support.....	13
<b>6</b>	<b>Processor Performance Monitoring .....</b>	<b>15</b>
6.1	Performance Monitor Programming Models.....	15
6.1.1	Workload Characterization .....	16
6.1.2	Profiling .....	19
6.1.3	Event Qualification .....	22
6.2	Performance Monitor State .....	27
6.2.1	Performance Monitor Control and Accessibility.....	27
6.2.2	Performance Counter Registers.....	30
6.2.3	Performance Monitor Overflow Status Registers (PMC[0,1,2,3]) .....	31

6.2.4	Intel® Itanium™ Instruction Address Range Check Register (PMC[13]) .....	32
6.2.5	Intel® Itanium™ Opcode Match Registers (PMC[8,9]).....	34
6.2.6	Intel® Itanium™ Data Address Range Check (PMC[11]).....	35
6.2.7	Event Address Registers (PMC[10,11]/PMD[0,1,2,3,17]) .....	36
6.2.8	Intel® Itanium™ Branch Trace Buffer .....	42
6.2.9	Processor Reset, PAL Calls, and Low Power State.....	46
6.2.10	References .....	47
<b>7</b>	<b>Performance Monitor Events .....</b>	<b>49</b>
7.1	Categorization of Events .....	49
7.2	Basic Events.....	49
7.3	Instruction Execution .....	50
7.4	Cycle Accounting Events.....	52
7.5	Branch Events.....	54
7.6	Memory Hierarchy .....	56
7.6.1	L1 Instruction Cache and Prefetch .....	59
7.6.2	L1 Data Cache .....	60
7.6.3	L2 Unified Cache.....	62
7.6.4	L3 Unified Cache.....	62
7.6.5	Frontside Bus .....	63
7.7	System Events .....	65
7.8	Performance Monitor Event List.....	67
<b>8</b>	<b>Model Specific Behavior for IA-32 Instruction Execution .....</b>	<b>113</b>
8.1	Processor Reset and Initialization .....	113
8.2	New JMPE Instruction .....	113
8.3	System Management Mode (SMM).....	114
8.4	CPUID Instruction Return Values for Caches and TLBs of the Intel® Itanium™ Processor .....	114
8.5	Machine Check Abort (MCA).....	115
8.6	Model Specific Registers.....	115
8.7	Cache Modes .....	115
8.8	10-byte Floating-point Operand Reads and Writes .....	115
8.9	Floating-point Data Segment State .....	116
8.10	Writes to Reserved Bits during FXSAVE.....	116
8.11	Setting the Access/Dirty (A/D) Bit on Accesses that Cross a Page Boundary.....	116
8.12	Enhanced Floating-point Instruction Accuracy .....	116
8.13	RCPSS, RCPPS, RSQRTSS, RSQRTPS Instruction Differences .....	117
8.14	Read/Write Access Ordering.....	117
8.15	Multiple IOAPIC Redirection Table Entries .....	117
8.16	Self Modifying Code (SMC).....	117
8.17	Raising an Alignment Check (AC) Fault.....	118
8.18	Maximum Number of Processors Supported in MP System Running Legacy IA-32 OS (IA-32 System Environment).....	118



## Figures

6-1	Time-based Sampling .....	16
6-2	Cycle Accounting in the Intel® Itanium™ Architecture.....	18
6-3	Event Histogram by Program Counter .....	20
6-4	Intel® Itanium™ Processor Event Qualification .....	23
6-5	Instruction Tagging Mechanism in the Intel® Itanium™ Processor.....	24
6-6	Single Process Monitor .....	26
6-7	Multiple Process Monitor .....	26
6-8	System Wide Monitor .....	27
6-9	Intel® Itanium™ Processor Performance Monitor Register Model.....	28
6-10	Processor Status Register (PSR) Fields for Performance Monitoring.....	29
6-11	Performance Monitor PMC Register Control Fields (PMC[4,5,6,7,10,11,12]) .....	29
6-12	Intel® Itanium™ Processor Generic PMD Registers (PMD[4,5,6,7]) .....	30
6-13	Intel® Itanium™ Processor Generic PMC Registers (PMC[4,5]) .....	30
6-14	Intel® Itanium™ Processor Generic PMC Registers (PMC[6,7]) .....	30
6-15	Intel® Itanium™ Processor Performance Monitor Overflow Status Registers (PMC[0,1,2,3]) .....	32
6-16	Intel® Itanium™ Processor Instruction Address Range Check Register (PMC[13]) .....	32
6-17	Opcode Match Registers (PMC[8,9]) .....	34
6-18	Instruction Event Address Configuration Register (PMC[10]) .....	36
6-19	Instruction Event Address Register Format (PMD[0,1]) .....	37
6-20	Data Event Address Configuration Register (PMC[11]) .....	39
6-21	Data Event Address Register Format (PMD[2,3,17]) .....	40
6-22	Intel® Itanium™ Branch Trace Buffer Configuration Register (PMC[12]) .....	42
6-23	Branch Trace Buffer Register Format (PMD[8-15]) .....	44
6-24	Intel® Itanium™ Branch Trace Buffer Index Register Format (PMD[16]) .....	46
7-1	Event Monitors in the Intel® Itanium™ Processor Memory Hierarchy .....	57
7-2	L1 Instruction Cache and Prefetch Monitors .....	59
7-3	L1 Data Cache Monitors.....	60
7-4	Instruction and Data TLB Monitors.....	67

## Tables

4-1	Intel® Itanium™ Processor WCB Eviction Conditions .....	9
6-1	Average Latency per Request and Requests per Cycle Calculation Example.....	18
6-2	Intel® Itanium™ Processor EARs and Branch Trace Buffer .....	21
6-3	Intel® Itanium™ Processor Event Qualification Modes.....	24
6-4	Intel® Itanium™ Processor Performance Monitor Register Set .....	29
6-5	Intel® Itanium™ Processor Generic PMD Register Fields .....	30
6-6	Intel® Itanium™ Processor Generic PMC Register Fields (PMC[4,5,6,7]) .....	31
6-7	Intel® Itanium™ Processor Performance Monitor Overflow Register Fields (PMC[0,1,2,3]) .....	32
6-8	Intel® Itanium™ Processor Instruction Address Range Check Register Fields (PMC[13]) .....	33
6-9	Intel® Itanium™ Processor Instruction Address Range Check by Instruction Set .....	33
6-10	Opcode Match Register Fields (PMC[8,9]).....	34
6-11	Instruction Event Address Configuration Register Fields (PMC[10]).....	37
6-12	Instruction EAR (PMC[10]) umask Field in Cache Mode (PMC[10].tlb=0) .....	38
6-13	Instruction EAR (PMD[0,1]) in Cache Mode (PMC[10].tlb=0).....	38
6-14	Instruction EAR (PMC[10]) umask Field in TLB Mode (PMC[10].tlb=1) .....	38
6-15	Instruction EAR (PMD[0,1]) in TLB Mode (PMC[10].tlb=1) .....	39
6-16	Data Event Address Configuration Register Fields (PMC[11]).....	39
6-17	PMC[11] Mask Fields in Data Cache Load Miss Mode (PMC[11].tlb=0).....	40
6-18	PMD[2,3,17] Fields in Data Cache Load Miss Mode (PMC[11].tlb=0) .....	40
6-19	PMC[11] Unmask Field in TLB Miss Mode (PMC[11].tlb=1) .....	41
6-20	PMD[2,3,17] Fields in TLB Miss Mode (PMC[11].tlb=1).....	41
6-21	Intel® Itanium™ Branch Trace Buffer Configuration Register Fields (PMC[12]).....	43
6-22	Intel® Itanium™ Branch Trace Buffer Register Fields (PMD[8-15]).....	45
6-23	Intel® Itanium™ Branch Trace Buffer Index Register Fields (PMD[16]) .....	46
6-24	Information Returned by PAL_PERF_MON_INFO for the Intel® Itanium™ Processor .....	46
7-1	Intel® Itanium™ and IA-32 Instruction Set Execution and Retirement Monitors.....	50
7-2	Intel® Itanium™ and IA-32 Instruction Set Execution and Retirement Performance Metrics .....	50
7-3	Instruction Issue and Retirement Events.....	50
7-4	Instruction Issue and Retirement Events (Derived).....	50
7-5	Floating-point Execution Monitors .....	51
7-6	Floating-point Execution Monitors (Derived) .....	51
7-7	Control and Data Speculation Monitors.....	51
7-8	Control/Data Speculation Performance Metrics .....	52
7-9	Memory Events .....	52
7-10	Stall Cycle Monitors.....	53
7-11	Stall Cycle Monitors (Derived).....	54
7-12	Branch Monitors .....	54
7-13	Branch Monitors (Derived) .....	55
7-14	Derived Memory Hierarchy Monitors.....	58
7-15	Cache Performance Ratios .....	58
7-16	L1 Instruction Cache and Instruction Prefetch Monitors.....	59
7-17	L1 Data Cache Monitors.....	60
7-18	L2 Cache Monitors .....	62
7-19	L3 Cache Monitors .....	62
7-20	Bus Events .....	63



7-21	Frontside Bus Monitors (Derived).....	64
7-22	Unit Masks for Qualifying Bus Transaction Events by Initiator .....	64
7-23	Conventions for Frontside Bus Transactions.....	64
7-24	Bus Events by Snoop Response .....	65
7-25	Bus Performance Metrics .....	65
7-26	System and TLB Monitors .....	66
7-27	System and TLB Monitors (Derived) .....	66
7-28	TLB Performance Metrics.....	67
7-29	Slot unit mask for BRANCH_TAKEN_SLOT .....	83
7-30	Retired Event Selection by Opcode Match.....	96
7-31	Unit Mask Bits {19:16} for L2_FLUSH_DETAILS Event .....	101
7-32	Unit Mask Bits {19:18} for PIPELINE_FLUSH Event.....	109
8-1	Encoding of Cache and TLB Return Values for the Intel® Itanium™ Processor.....	114
8-2	EAX, EBX, ECX, and EDX Return Values for the Intel® Itanium™ Processor .....	115





The *Intel® Itanium™ Processor Reference Manual for Software Development* describes model-specific architectural features incorporated into the Intel® Itanium™ processor, the first processor based on the Itanium architecture. This document (Document number 245320) has been re-titled. In previous revisions, it was titled the *Intel® Itanium™ Architecture Software Developer's Manual, Volume 4: Itanium™ Processor Programmer's Guide*.

## 1.1 Overview of the Intel® Itanium™ Processor Reference Manual for Software Development

Chapter 1, “About this Manual” provides an overview of four volumes in the *Intel® Itanium Architecture Software Developer's Manual*.

Chapter 2, “Register Stack Engine Support” summarizes Register Stack Engine (RSE) support provided by the Itanium processor.

Chapter 3, “Virtual Memory Management Support” details size of physical and virtual address, region register ID, and protection key register implemented on the Itanium processor.

Chapter 4, “Processor Specific Write Coalescing (WC) Behavior” describes the behavior of write coalesce (also known as Write Combine) on the Itanium processor.

Chapter 5, “Model Specific Instruction Implementation” describes model specific behavior of Itanium instructions on the Itanium processor.

Chapter 6, “Processor Performance Monitoring” defines the performance monitoring features which are specific to the Itanium processor. This chapter outlines the targeted performance monitor usage models and describes the Itanium processor specific performance monitoring state.

Chapter 7, “Performance Monitor Events” summarizes the Itanium processor events and describes how to compute commonly used performance metrics for Itanium processor events.

Chapter 8, “Model Specific Behavior for IA-32 Instruction Execution” describes some of the key differences between an Itanium processor executing IA-32 instructions and the Intel® Pentium® III processor.

## 1.2 Terminology

The following definitions are for terms related to the Itanium architecture and will be used throughout this document:

**Instruction Set Architecture (ISA)** – Defines application and system level resources. These resources include instructions and registers.

**Itanium Architecture** – The new ISA with 64-bit instruction capabilities, new performance-enhancing features, and support for the IA-32 instruction set.

**IA-32 Architecture** – The 32-bit and 16-bit Intel Architecture as described in the *Intel Architecture Software Developer's Manual*.

**Itanium System Environment** – The operating system environment that supports the execution of both IA-32 and Itanium-based code.

**IA-32 System Environment** – The operating system privileged environment and resources as defined by the *Intel Architecture Software Developer's Manual*. Resources include virtual paging, control registers, debugging, performance monitoring, machine checks, and the set of privileged instructions.

**Itanium-based Firmware** – The Processor Abstraction Layer (PAL) and System Abstraction Layer (SAL).

**Processor Abstraction Layer (PAL)** – The firmware layer which abstracts processor features that are implementation dependent.

**System Abstraction Layer (SAL)** – The firmware layer which abstracts system features that are implementation dependent.

## 1.3 Related Documents

The following documents can be downloaded at the Intel's Developer Site at <http://developer.intel.com>:

- **Intel® Itanium™ Processor Reference Manual for Software Development** – This document (Document number 245320) describes model-specific architectural features incorporated into the Intel® Itanium™ processor, the first processor based on the Itanium architecture. This document has been re-titled and replaces the *Intel® Itanium™ Architecture Software Developer's Manual, Volume 4: Itanium™ Processor Programmer's Guide*.
- **Intel® Architecture Software Developer's Manual** – This set of manuals describes the Intel 32-bit architecture. They are readily available from the Intel Literature Department by calling 1-800-548-4725 and requesting Document Numbers 243190, 243191 and 243192.
- **Itanium™ Software Conventions and Runtime Architecture Guide** – This document (Document number 245358) defines general information necessary to compile, link, and execute a program on an Itanium-based operating system.
- **Itanium™ Processor Family System Abstraction Layer Specification** – This document (Document number 245359) specifies requirements to develop platform firmware for Itanium-based systems.
- **Extensible Firmware Interface Specification** – This document defines a new model for the interface between operating systems and platform firmware.

## 1.4 Revision History

Date of Revision	Revision Number	Description
December 2001	2.0	<p>Initial release of re-titled document. This document (Document number 245320) replaces the former <i>Intel® Itanium™ Architecture Software Developer's Manual, Volume 4: Itanium™ Processor Programmer's Guide</i>.</p> <p>Performance monitoring changes (Section 7.8).</p> <p>Revised Chapter 7 Performance Monitoring Events (new Section 7.6.5, Frontside Bus; added bus monitors to Section 7.8, Event List; misc. changes and fixes).</p> <p>Revised IBR and DBR addressing (Section 6.2.4).</p> <p>IA-32 related changes (Section 8).</p> <p>Miscellaneous performance monitoring events changes (Chapter 7).</p>
July 2000	1.1	<p>Reformatted the Performance Monitor Events chapter for readability and ease of use (no changes to any of the events except for renaming of some); events are listed in alphabetical order (Chapter 7).</p>
January 2000	1.0	<p>Initial release of document.</p>



## 2.1 RSE Modes

The Itanium processor implements the enforced lazy RSE mode. Refer to [Chapter 6, “Register Stack Engine”](#) in [Volume 2](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a description of the RSE modes.

## 2.2 RSE and Clean Register Stack Partitions

On the Itanium processor, the internal RSE pointer `RSE.BSPLoad` is always equal to `AR.BSPStore`, meaning that the size of the clean register stack partition is always zero. This implies that, on the Itanium processor, a `flushrs` instruction will create a dirty region of size zero and an invalid region of size equal to `96 - CFM.sof`. On other implementations that maintain a clean partition, `flushrs` behavior may differ by creating a clean register stack partition in addition to an invalid partition and a zero-sized dirty partition. As a result, the Itanium processor's RSE may perform more mandatory fills upon a branch-return (`br.ret`) or `rfi` following a `flushrs` instruction than an implementation that maintains a clean partition.



# **Virtual Memory Management Support 3**

---

## **3.1 Page Size Supported**

The following page sizes are supported on the Itanium processor: 4K, 8K, 16K, 64K, 256K, 1M, 4M, 16M and 256M bytes.

## **3.2 Physical and Virtual Addresses**

The Itanium architecture requires that a processor implement at least 54 virtual address bits and 32 physical address bits. The Itanium processor implements 54 virtual address bits (51 address bits plus 3 region index bits) and 44 physical address bits.

## **3.3 Region Register ID**

The Itanium processor implements the minimum region register IDs allowed by the Itanium architecture. The region register ID contains 18 bits.

## **3.4 Protection Key Register**

The Itanium architecture requires a minimum of 16 protection key registers, each at least as wide as the region register IDs. The Itanium processor implements 16 protection key registers, each 21 bits wide.





# Processor Specific Write Coalescing (WC) Behavior

## 4.1 Write Coalescing

For increased performance of uncacheable references to frame buffers, previous Intel IA-32 processors defined the Write Coalescing (WC) memory type. WC coalesces streams of data writes into a single larger bus write transaction. Refer to the *Intel® Architecture Software Developer's Manual* for additional information.

On the Itanium processor, WC loads are performed directly from memory and not from coalescing buffers. It has a separate 2-entry, 64-byte Write Coalesce Buffer (WCB) which is used exclusively for WC accesses. Each byte in the line has a valid bit. If all the valid bits are true, then the line is full and will be evicted (or flushed) by the processor.

**Note:** WC behavior of the Itanium processor in the IA-32 System Environment is similar to the Pentium III processor. Refer to the *Intel® Architecture Software Developer's Manual* for more information.

## 4.2 WC Buffer Eviction Conditions

To ensure consistency with memory, the WCB is flushed on the following conditions (both entries are flushed). These conditions are followed when the processor is operating in the Itanium System Environment:

**Table 4-1. Intel® Itanium™ Processor WCB Eviction Conditions**

Eviction Condition	Intel® Itanium™ Instructions
Memory Fence (mf)	mf
<b>Architectural Conditions for WCB Flush</b>	
Memory Release ordering (op.rel)	st.rel, cmpxchg.rel, fetchadd.rel, ptc.g
Flush Cache (fc) hit on WCB	yes
Flush Write Buffers (fwb)	yes
Any UC load	no <sup>a</sup>
Any UC store	no <sup>a</sup>
UC load or ifetch hits WCB	no <sup>a</sup>
UC store hits WCB	no <sup>a</sup>
WC load/ifetch hits WCB	
WC store hits WCB	

a. Itanium architecture doesn't require the WC buffers to be coherent w.r.t to UC load/store operations.

## 4.3 WC Buffer Flushing Behavior

As mentioned previously, the Itanium processor WCB contains two entries. The WC entries are flushed in the same order as they are allocated. That is, the entries are flushed in written order. This flushing order applies only to a “well-behaved” stream. A “well-behaved” stream writes one WC entry at a time and does not write the second WC entry until the first one is full.

In the absence of platform retry or deferral, the flushing rule implies that the WCB entries are always flushed in a program written order for a “well-behaved” stream, even in the presence of interrupts. For example, consider the following scenario: if software issues a “well-behaved” stream, but is interrupted in the middle; one of the WC entries could be partially filled. The WCB (including the partially filled entry) could be flushed by the OS kernel code or by other processes. When the interrupted context resumes, it sends out the remaining line and then moves on to fill the other entry. Note that the resumed context could be interrupted again in the middle of filling up the other entry, causing both entries to be partially filled when the interrupt occurs.

For streams that do not conform to the above “well-behaved” rule, the order in which the WC buffer is flushed is random.

WCB eviction is performed for full lines by a single 64-byte bus transaction in a stream of 8-byte packages. For partially full lines, the WCB is evicted using up to eight 8-byte transactions with the proper byte enables. When flushing, WC transactions are given the highest priority of all external bus operations.

# Model Specific Instruction Implementation

## 5

This section describes how Itanium instructions with processor implementation-specific features, behave on the Intel Itanium processor.

### 5.1 `ld.bias`

If the instruction hits L1D<sup>1</sup> or L2 cache and the state of the line is exclusive (E) or modified (M), the line is returned and remains in cache; no external bus traffic is generated. If the line is shared (S) or invalid (I) or the instruction misses the L2, it is treated as a store miss by the L3/bus. The line is returned and stored in E state by the processor in the L2 and L3 cache.

Please refer to [page 3:135](#) in [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a detailed description of the `ld` instruction.

### 5.2 `lfetch Exclusive Hint`

The exclusive hint in the `lfetch` instruction allows the cache line to be fetched in an exclusive (E) state. On the Itanium processor, an `lfetch` transaction that has a snoop hit will be cached in a shared (S) state; otherwise, it is cached in an exclusive state.

Please refer to [page 3:146](#) in [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a detailed description of the `lfetch` instruction.

### 5.3 `fwb`

The Itanium processor implements the flush write-back buffer (`fwb`) instruction. This instruction carries a weak memory attribute and causes the coalescing buffer to be flushed. The L1D and L2 store buffers are not flushed.

Please refer to [page 3:126](#) in [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a detailed description of the `fwb` instruction.

---

1. The Itanium processor cache hierarchy consists of the following levels: on-chip L1I, L1D, L2 caches, and off-chip L3 cache.

## 5.4 thash

The Itanium architecture defines a `thash` instruction for generating the hash address for long format VHPT. `thash` is implementation specific. On the Itanium processor, since the hashing function is performed in the HPW, the HPW will generate the VHPT Entry which corresponds to the virtual address supplied. The hashing function is given in the following pseudo-code:

```

If (GR[r3].nat = '1 or unimplemented virtual address bits) then {
GR[r1] = '0 ;                               // treated as a speculative access.
GR[r1].nat = '1;
}
else {
Mask = (2^PTA.size) - 1;
HPN = VA{50:0} >> RR[VA{63:61}].ps;      // Hash Page Number unsigned right shift.
// mov 2 RR checks for supported ps
if (PTA.vf=32) {                            // 32B PTE (Long format)
Hash_Index = HPN ^ (zero{63:18} || rid{17:0})
VHPT_Offset = Hash_Index << 5 ;
}
if (PTA.vf=8) {                             // 8B PTE
Hash_Index = HPN ;
VHPT_Offset = Hash_Index << 3;
}
GR[r1] = (PTA.base{63:61} << 61)
|| ((PTA.base{60:15} & ~Mask{60:15}) ||
(VHPT_Offset{60:15} & Mask{60:15})) << 15)
|| VHPT_Offset{14:0} ;
}
}

```

Please refer to [page 3:234](#) in [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a detailed description of the `thash` instruction.

## 5.5 ttag

The Itanium architecture defines the `ttag` instruction for generating the tag for a long format VHPT entry. `ttag` is implementation specific. The HPW will generate the tag for the long format VHPT entry which corresponds to the virtual address supplied. The function is:

```

If (GR[r3].nat = '1 or unimplemented virtual address bits) then {
GR[r1] = '0 ;
GR[r1].nat = '1;
}
else {
GR[r1] = (VA{50:0} >> RR[VA{63:61}].PS) ^
((zero{5:0} || RR[VA{63:61}].RID{17:0}) << 39);
}
}

```

Please refer to [page 3:238](#) in [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a detailed description of the `ttag` instruction.

## 5.6 **ptc.e**

On the Itanium processor, a single `ptc.e` purges all translation cache (TC) entries in both the instruction and data TLBs. The caches are not flushed.

Please refer to [page 3:202](#) in [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a detailed description of the `ptc` instruction.

## 5.7 **mf.a**

In the Itanium architecture, the `mf.a` instruction is a memory acceptance fence for UC transactions only. On the Itanium processor, `mf.a` is implemented as an acceptance fence for both cacheable and UC data transactions (but not I fetches). The processor stalls until all data buffers in the L2 and bus are empty. This does not include buffers for instruction and L3 WB buffer in the bus request queue.

Please refer to [page 3:149](#) in [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for a detailed description of the `mf` instruction.

## 5.8 **Prefetch Behavior**

The Itanium processor does not initiate prefetches with post-increment loads.

## 5.9 **Temporal and Non-temporal Hints Support**

Itanium architecture provides memory locality hints for data accesses that can be used for allocation control in the processor cache hierarchy. For more details on this topic, please refer to [Volume 1](#) of the *Intel® Itanium™ Architecture Software Developer's Manual*, [Section 4.4.6](#). Implementation of locality hints is left as an implementation-specific feature on processors based on the Itanium architecture.

On the Itanium processor, four types of memory locality hints are implemented: `t1`, `nt1`, `nt2` and `nta`. The Itanium processor does not support a non-temporal buffer; instead, non-temporal accesses are allocated in L2 cache with biased replacement.



# Processor Performance Monitoring 6

---

This chapter defines the performance monitoring features on the Itanium processor. The Itanium processor provides four 32-bit performance counters, more than 50 monitorable events, and several advanced monitoring capabilities. This chapter outlines the targeted performance monitor usage models, defines the software interface and programming model, and lists the set of monitored events.

Itanium architecture incorporates architected mechanisms that allow software to actively and directly manage performance critical processor resources such as branch prediction structures, processor data and instruction caches, virtual memory translation structures, and more. To achieve the highest performance levels, dynamic processor behavior can be monitored and fed back into the code generation process to improve observed run-time behavior or to expose higher levels of instruction level parallelism. One can quantify and measure behavior of real-world Itanium-based applications, tools and operating systems. These measurements will be critical for compiler optimizations and the efficient use of several architectural features such as speculation, predication, and more.

The remainder of this chapter is split into the following two subsections:

- [Section 6.1, "Performance Monitor Programming Models"](#) discusses how performance monitors are used and presents various Itanium processor performance monitoring programming models.
- [Section 6.2, "Performance Monitor State"](#) defines the Itanium processor specific PMC/PMD performance monitoring registers.

## 6.1 Performance Monitor Programming Models

This section introduces the Itanium processor performance monitoring features from a programming model point-of-view and describes how the different event monitoring mechanisms can be used effectively. The Itanium processor performance monitor architecture focuses on the following two usage models:

- **Workload Characterization:** the first step in any performance analysis is to understand the performance characteristics of the workload under study. [Section 6.1.1, "Workload Characterization"](#) discusses the Itanium processor support for workload characterization.
- **Profiling:** profiling is used by application developers and profile-guided compilers. Application developers are interested in identifying performance bottlenecks and relating them back to their code. Their primary objective is to understand which program location caused performance degradation at the module, function, and basic block level. For optimization of data placement and the analysis of critical loops, instruction level granularity is desirable. Profile-guided compilers that use advanced Itanium architectural features such as predication and speculation benefit from run-time profile information to optimize instruction schedules. The Itanium processor supports instruction granular statistical profiling of branch mispredicts and cache misses. Details of the Itanium processor's profiling support are described in [Section 6.1.2, "Profiling"](#).

Whenever monitoring overhead is irrelevant, but accuracy is the primary objective, system and processor designers may resort to tracing processor activity at the system or the processor bus interface. However, trace based performance analysis and hardware tracing of the Itanium processor are beyond the scope of this documentation.

## 6.1.1 Workload Characterization

The first step in any performance analysis is to understand the performance characteristics of the workload under study. There are two fundamental measures of interest: event rates and program cycle break down.

- **Event Rate Monitoring:** Event rates of interest include average retired instructions-per-clock (IPC), data and instruction cache miss rates, or branch mispredict rates measured across the entire application. Characterization of operating systems or large commercial workloads (e.g. OLTP analysis) requires a system-level view of performance relevant events such as TLB miss rates, VHPT walks/second, interrupts/second or bus utilization rates. [Section 6.1.1.1, "Event Rate Monitoring"](#) discusses event rate monitoring.
- **Cycle Accounting:** The cycle break-down of a workload attributes a reason to every cycle spent by a program. Apart from a program's inherent execution latency, extra cycles are usually due to pipeline stalls and flushes. [Section 6.1.1.4, "Cycle Accounting"](#) discusses cycle accounting.


### 6.1.1.1 Event Rate Monitoring

Event rate monitoring determines event rates by reading processor event occurrence counters before and after the workload is run and then computing the desired rates. For instance, two basic Itanium processor events that count the number of retired Itanium instructions (IA64\_INST\_RETIRED) and the number of elapsed clock cycles (CPU\_CYCLES) allow a workload's instructions per cycle (IPC) to be computed as follows:

$$IPC = (IA64\_INST\_RETIRED_{t1} - IA64\_INST\_RETIRED_{t0}) / (CPU\_CYCLES_{t1} - CPU\_CYCLES_{t0})$$

Time-based sampling is the basis for many performance debugging tools [VTune™, gprof, Windows NT\*]. As shown in [Figure 6-1](#), time-based sampling can be used to plot the event rates over time, and can provide insights into the different phases the workload moves through.

**Figure 6-1. Time-based Sampling**





On the Itanium processor, many event types (e.g. TLB misses or branch mispredicts) are limited to a rate of one per clock cycle. These are referred to as “single occurrence” events. However, in the Itanium processor multiple events of the same type may occur in the same clock. We refer to such events as “multi-occurrence” events. An example of a multi-occurrence events on the Itanium processor is data cache misses (up to two per clock). Multi-occurrence events, such as the number of entries in the memory request queue, can be used to derive average number and average latency of memory accesses. The next two sections describe the basic Itanium processor mechanisms for monitoring single and multi-occurrence events.

### 6.1.1.2 Single Occurrence Events and Duration Counts

A single occurrence event can be monitored by any of the Itanium processor performance counters. For all single occurrence events a counter is incremented by up to one per clock cycle. Duration counters that count the number of clock cycles during which a condition persists are considered “single occurrence” events. Examples of single occurrence events on the Itanium processor are TLB misses, branch mispredictions, or cycle-based metrics.

### 6.1.1.3 Multi-occurrence Events, Thresholding and Averaging

Events that, due to hardware parallelism, may occur at rates greater than one per clock cycle are termed “multi-occurrence” events. Examples of such events on the Itanium processor are retired instructions or the number of live entries in the memory request queue. The Itanium processor’s four performance counters are asymmetrical. While all counters handle single-occurrence and multi-occurrence events with event rates up to three per cycle, only two counters can handle multi-occurrence events with event rates up to seven per cycle. For details, see [Section 6.2.2, "Performance Counter Registers"](#).

Thresholding capabilities are available in the Itanium processor’s multi-occurrence counters and can be used to plot an event distribution histogram. When a non-zero threshold is specified, the monitor is incremented by one in every cycle in which the observed event count exceeds that programmed threshold. This allows questions such as “for how many cycles did the memory request queue contain more than two entries?” or “during how many cycles did the machine retire more than three instructions?” to be answered. This capability allows micro-architectural buffer sizing experiments to be supported by real measurements. By running a benchmark with different threshold values, a histogram can be drawn up that may help to identify the performance “knee” at a certain buffer size.

For overlapping concurrent events, such as pending memory operations, the average number of concurrently outstanding requests and the average number of cycles that requests were pending is of interest. To calculate the average number or latency of multiple outstanding requests in the memory queue, we need to know the total number of requests ( $n_{total}$ ) and, in each cycle, the number of live requests per cycle ( $n_{live}/cycle$ ). By summing up the live requests ( $n_{live}/cycle$ ) using a multi-occurrence counter  $\Sigma n_{live}$  is directly measured by hardware. We can now calculate the average number of requests and the average latency as follows:

- Average outstanding requests/cycle =  $\Sigma n_{live} / \Delta t$
- Average latency per request =  $\Sigma n_{live} / n_{total}$

An example of this calculation is given in [Table 6-1](#), in which the average outstanding requests/cycle =  $15/8 = 1.825$ , and the average latency per request =  $15/5 = 3$  cycles.

**Table 6-1. Average Latency per Request and Requests per Cycle Calculation Example**

<b>Time [Cycles]</b>	1	2	3	4	5	6	7	8
<b># Requests In</b>	1	1	1	1	1	0	0	0
<b># Requests Out</b>	0	0	0	1	1	1	1	1
<b><math>n_{live}</math></b>	1	2	3	3	3	2	1	0
<b><math>\Sigma n_{live}</math></b>	1	3	6	9	12	14	15	15
<b><math>n_{total}</math></b>	1	2	3	4	5	5	5	5

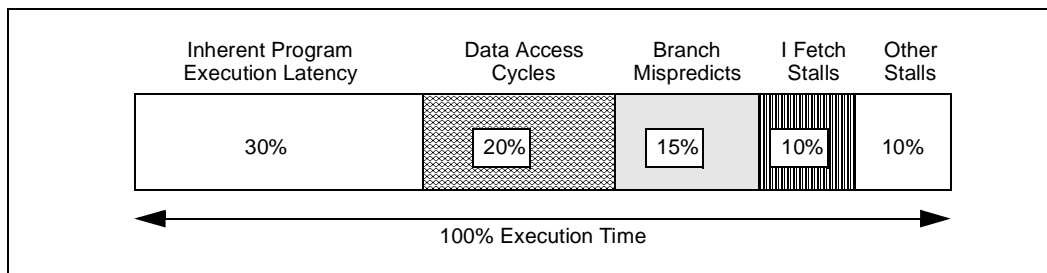
The Itanium processor provides the following capabilities to support event rate monitoring:

- Clock cycle counter
- Retired instruction counter
- Event occurrence and duration counters
- Multi-occurrence counters with thresholding capability

### 6.1.1.4 Cycle Accounting

While event rate monitoring counts the number of events, it does not tell us whether the observed events are contributing to a performance problem. A commonly used strategy is to plot multiple event rates and correlate them with the measured instructions per cycle (IPC) rate. If a low IPC occurs concurrently with a peak of cache miss activity, chances are that cache misses are causing a performance problem. To eliminate such guess work, the Itanium processor provides a set of cycle accounting monitors based on the Itanium architecture, that break-down the number of cycles that are lost due to various kinds of micro-architectural events. As shown in Figure 6-2, this lets us account for every cycle spent by a program and therefore provides insight into an application's micro-architectural behavior. Note that cycle accounting is different from simple stall or flush duration counting. Cycle accounting is based on the machine's actual stall and flush conditions and accounts for overlapped pipeline delays, while simple stall or flush duration counters do not. Cycle accounting determines a program's cycle break-down by stall and flush reasons, while simple duration counters are useful in determining cumulative stall or flush latencies.

**Figure 6-2. Cycle Accounting in the Intel® Itanium™ Architecture**



The Itanium processor cycle accounting monitors account for all major single and multi-cycle stall and flush conditions. Overlapping stall and flush conditions are prioritized in reverse pipeline order (i.e. delays that occur later in the pipe and that overlap with earlier stage delays are reported as being caused later in the pipeline). The eight stall and flush reasons are prioritized in the following order:

1. Back-end Flush Cycles: cycles lost due to branch mispredictions, ALAT flushes, serialization flushes, failed control speculation flushes, MMU-IEU bypasses and other exceptions.

2. Data Access Cycles: cycles lost when instructions stall waiting for their source operands from the memory subsystem, and when memory flushes arise (L1D way mispredictions, DTC flushes).
3. Scoreboard Dependency Cycles: cycles lost when instructions stall waiting for their source operands from non-load instructions; this includes FP-related flushes.
4. RSE Active Cycles: stalls due to register stack spills to and fills from the backing store in memory.
5. Issue Limit Cycles: dispersal breaks due to stops, port over-subscription or asymmetries.
6. Instruction Access Cycles: instruction fetch stalls due to L1I or ITLB misses.
7. Taken Branch Cycles: bubbles incurred on correct taken branch predictions.

Four of the eight categories (1, 2, 3, 6) are directly measurable as Itanium processor events. The other four categories (4, 5, 7, 8) are not measured directly. Instead, four combined categories are available as the Itanium processor events: pipeline flush cycles (1+7), memory cycles (2+4), dependency cycles (3+5), and unstalled back-end cycles (6+8). For details refer to [Section 7.4, “Cycle Accounting Events” on page 50](#).

## 6.1.2 Profiling

Profiling is used by application developers and profile-guided compilers, optimizing linkers and run-time systems. Application developers are interested in identifying performance bottlenecks and relating them back to their source code. Based on profile feedback developers can make changes to the high-level algorithms and data structures of the program. Compilers can use profile feedback to optimize instruction schedules by employing advanced Itanium architectural features such as predication and speculation.

To support profiling, performance monitor counts have to be associated with program locations. The following mechanisms are supported directly by the Itanium processor’s performance monitors:


- Program Counter Sampling
- Miss Event Address Sampling: Itanium processor Event Address Registers (EARs) provide sub-pipeline length event resolution for performance critical events (instruction and data caches, branch mispredicts, instruction and data TLBs).
- Event Qualification: constrains event monitoring to a specific instruction address range, to certain opcodes or privilege levels.

These profiling features are presented in the next three subsections.

### 6.1.2.1 Program Counter Sampling

Application tuning tools like [VTune, gprof] use time-based or event-based sampling of the program counter and other event counters to identify performance critical functions and basic blocks. As shown in [Figure 6-3](#), the sampled points can be histogrammed by instruction addresses. For application tuning, statistical sampling techniques have been very successful, because the programmer can rapidly identify code hot-spots in which the program spends a significant fraction of its time or where certain event counts are high.

**Figure 6-3. Event Histogram by Program Counter**



Program counter sampling points the performance analysts at code hot-spots, but does not indicate what caused the performance problem. Inspection and manual analysis of the hot-spot region along with a fair amount of guess work are required to identify the root cause of the performance problem. On the Itanium processor, the cycle accounting mechanism (described in [Section 6.1.1.4, "Cycle Accounting"](#)) can be used to directly measure an application's micro-architectural behavior.

The Itanium architectural interval timer facilities (ITC and ITM registers) can be used for time-based program counter sampling. Event-based program counter sampling is supported by a dedicated performance monitor overflow interrupt mechanism described in detail in [Volume 2, Section 7.2.2, "Performance Monitor Overflow Status Registers \(PMC\[0\]..PMC\[3\]\)"](#).

To support program counter sampling, the Itanium processor provides the following mechanisms:

- Timer interrupt for time-based program counter sampling.
- Event count overflow interrupt for event-based program counter sampling.
- Hardware supported cycle accounting.

### 6.1.2.2 Miss Event Address Sampling

Program counter sampling and cycle accounting provide an accurate picture of cumulative micro-architectural behavior, but they do not provide the application developer with pointers to specific program elements (code locations and data structures) that repeatedly cause micro-architectural "miss events". In a cache study of the SPEC92 benchmarks, [Lebeck] used (trace based) cache miss profiling to gain performance improvements of 1.02 to 3.46 on various benchmarks by making simple changes to the source code. This type of analysis requires identification of instruction and data addresses related to micro-architectural "miss events" such as cache misses, branch mispredicts, or TLB misses. Using symbol tables or compiler annotations these addresses can be mapped back to critical source code elements. Like Lebeck, most performance analysts in the past have had to capture hardware traces and resort to trace driven simulation.

Due to the super-scalar issue, deep pipelining, and out-of-order instruction completion of today's microarchitectures, the sampled program counter value may not be related to the instruction address that caused a miss event. On a Pentium processor pipeline, the sampled program counter may be off by 2 dynamic instructions from the instruction that caused the miss event. On a Pentium Pro processor, this distance increases to approximately 32 dynamic instructions. On the Itanium processor it is approximately 48 dynamic instructions. If program counter sampling is used for miss event address identification on the Itanium processor, a miss event might be associated with an

instruction almost five dynamic basic blocks away from where it actually occurred (assuming that 10% of all instructions are branches). Therefore, it is essential for hardware to precisely identify an event's address.

The Itanium processor provides a set of *event address registers* (EARs) that record the instruction and data addresses of data cache misses for loads, the instruction and data addresses of data TLB misses, the instruction addresses of instruction TLB and cache misses. A four deep *branch trace buffer* captures sequences of branch instructions. [Table 6-2](#) summarizes the capabilities offered by the EARs and branch trace buffer. Exposing miss event addresses to software allows them to be monitored either by sampling or by code instrumentation. This eliminates the need for trace generation to identify and solve performance problems and enables performance analysis by a much larger audience on unmodified hardware.

**Table 6-2. Intel® Itanium™ Processor EARs and Branch Trace Buffer**

Event Address Register	Triggers on	What is Recorded
Instruction Cache	Instruction fetches that miss the L1 instruction cache (demand fetches only)	Instruction Address Number of cycles fetch was in flight.
Instruction TLB (ITLB)	Instruction fetch missed ITLB (demand fetches only)	Instruction Address Who serviced TLB miss: VHPT or software.
Data Cache	Load instructions that miss L1 data cache	Instruction Address Data Address Number of cycles load was in flight.
Data TLB (DTLB)	Data references that miss L1 DTLB	Instruction Address Data Address Who serviced TLB miss: L2 DTLB, VHPT or software.
Branch Trace Buffer	Branch Outcomes	Branch Instruction Address Branch Target Instruction Address Mispredict status and reason

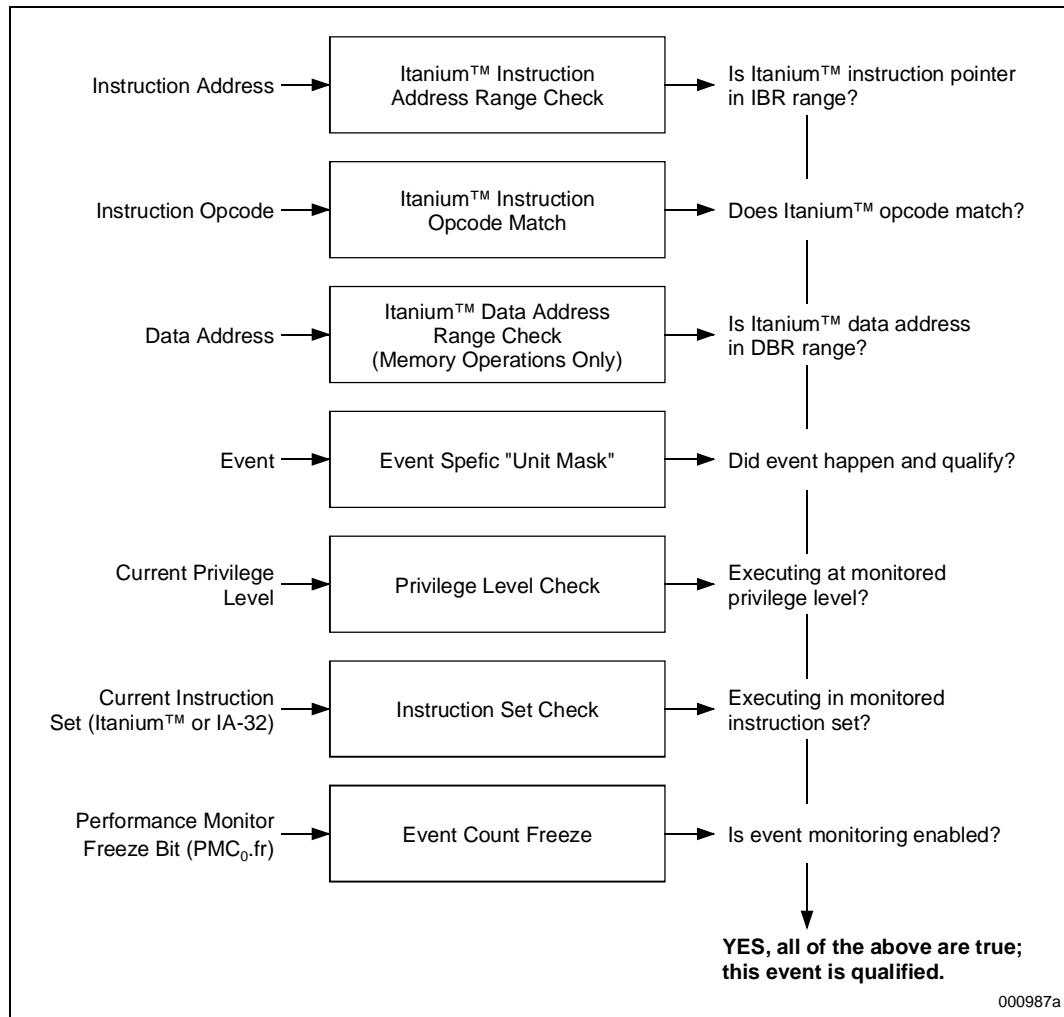
The Itanium processor EARs enable statistical sampling by configuring a performance counter to count, for instance, the number of data cache misses or retired instructions. The performance counter value is set up to interrupt the processor after a pre-determined number of events have been observed. The data cache event address register repeatedly captures the instruction and data addresses of actual data cache load misses. Whenever the counter overflows, miss event address collection is suspended until the event address register is read by software (this prevents software from capturing a miss event that might be caused by the monitoring software itself). When the counter overflows an interrupt is delivered to software, the observed event addresses are collected, and a new observation interval can be setup by rewriting the performance counter register. For time-based (rather than event-based) sampling methods, the event address registers indicate to software whether or not a qualified event was captured. Statistical sampling can achieve arbitrary event resolution by varying the number of events within an observation interval, and by increasing the number of observation intervals.

### 6.1.3 Event Qualification

On the Itanium processor, performance monitoring can be confined to a subset of all events. As shown in [Figure 6-4](#), events can be qualified for monitoring based on an instruction address range, a particular instruction opcode, a data address range, an event specific “unit-mask”, the privilege level and instruction set the event was caused by, and the status of the performance monitoring freeze bit (PMC[0].fr).

- **Itanium Instruction Address Range Check:** The Itanium processor allows event monitoring to be constrained to a programmable instruction address range. This enables monitoring of dynamically linked libraries (DLL), functions, or loops of interest in the context of a large Itanium-based application. The Itanium instruction address range check is applied at the instruction fetch stage of the pipeline and the resulting qualification is carried by the instruction throughout the pipeline. This enables conditional event counting at a level of granularity smaller than dynamic instruction length of the pipeline (approximately 48 instructions). The Itanium processor’s instruction address range check operates only during Itanium-based code execution (i.e. when PSR.is is zero). For details, see [Section 6.2.4, "Intel® Itanium™ Instruction Address Range Check Register \(PMC\[13\]\)"](#).
- **Itanium Instruction Opcode Match:** The Itanium processor provides two independent Itanium opcode match registers each of which match the currently issued instruction encodings with a programmable opcode match and mask function. The resulting match events can be selected as an event type for counting by the performance counters. This allows histogramming of instruction types, usage of destination and predicate registers as well as basic block profiling (through insertion of tagged nops). The opcode matcher operates only during Itanium-based code execution (i.e. when PSR.is is zero). Details are described in [Section 6.2.5, "Intel® Itanium™ Opcode Match Registers \(PMC\[8,9\]\)"](#).
- **Itanium Data Address Range Check:** The Itanium processor allows event collection for memory operations to be constrained to a programmable data address range. This enables selective monitoring of data cache miss behavior of specific data structures. For details, see [Section 6.2.6, "Intel® Itanium™ Data Address Range Check \(PMC\[11\]\)"](#).
- **Event Specific Unit Masks:** Some events allow the specification of “unit masks” to filter out interesting events directly at the monitored unit. For details, refer to the event pages in [Chapter 7, "Performance Monitor Events"](#).
- **Privilege Level:** Two bits in the processor status register are provided to enable selective process-based event monitoring. The Itanium processor supports conditional event counting based on the current privilege level; this allows performance monitoring software to break-down event counts into user and operating system contributions. For details on how to constrain monitoring by privilege level refer to [Section 6.2.1, "Performance Monitor Control and Accessibility"](#).
- **Instruction Set:** The Itanium processor supports conditional event counting based on the currently executing instruction set (Itanium architecture or IA-32) by providing two instruction set mask bits for each event monitor. This allows performance monitoring software to break-down event counts into Itanium-based and IA-32 contributions. For details, refer to [Section 6.2.1, "Performance Monitor Control and Accessibility"](#).
- **Performance Monitor Freeze:** Event counter overflows or software can freeze event monitoring. When frozen, no event monitoring takes place until software clears the monitoring freeze bit (PMC[0].fr). This ensures that the performance monitoring routines themselves, e.g. counter overflow interrupt handlers or performance monitoring context switch routines, do not “pollute” the event counts of the system under observation.

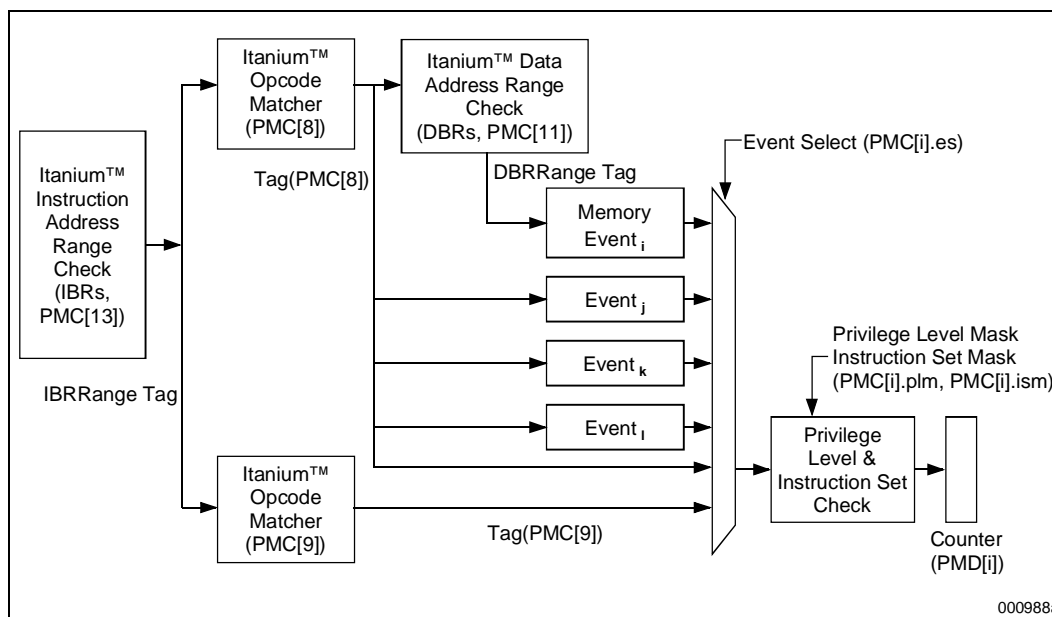
**Figure 6-4. Intel® Itanium™ Processor Event Qualification**



### 6.1.3.1 Combining Opcode Matching, Instruction, and Data Address Range Check

The Itanium processor allows various event qualification mechanisms to be combined by providing the instruction tagging mechanism shown in Figure 6-5. Instruction address range check and opcode matching are available only for Itanium-based code; they are disabled when IA-32 code is executing.

**Figure 6-5. Instruction Tagging Mechanism in the Intel® Itanium™ Processor**



During Itanium instruction execution (PSR.is is zero), the instruction address range check is applied first. The resulting address range check tag (IBRRangeTag) is passed to two opcode matchers that combine the instruction address range check with the opcode match. Each of the two combined tags (Tag(PMC[8]) and Tag(PMC[9])) can be counted as a retired instruction count event (for details refer to event description IA64\_TAGGED\_INST\_RETIRE in Table 7-3 “Instruction Issue and Retirement Events” on page 48).

One of the combined Itanium address range and opcode match tags, Tag(PMC[8]), qualifies all down-stream pipeline events. Events in the memory hierarchy (L1 and L2 data cache and data TLB events) can further be qualified using a data address DBRRangeTag.

As summarized in Table 6-3, data address range checking can be combined with opcode matching and instruction range checking on the Itanium processor. Additional event qualifications based on the current privilege level and the current instruction set can be applied to all events and are discussed in Section 6.1.3.2, “Privilege Level Constraints” and Section 6.1.3.3, “Instruction Set Constraints”.

**Table 6-3. Intel® Itanium™ Processor Event Qualification Modes**

Event Qualification Modes	Instr. Address Range Check PMC[13].ta	Opcode Matching PMC[8]	Data Address Range Check PMC[11].pt
Unconstrained Monitoring (all events)	1	0xffff_ffff_ffff_ffff	1
Instruction Address Range Check only	0	0xffff_ffff_ffff_ffff	1
Opcode Matching only	1	Desired Opcodes	1
Data Address Range Check only	1	0xffff_ffff_ffff_ffff	0
Instruction Address Range Check and Opcode Matching	0	Desired Opcodes	1



**Table 6-3. Intel® Itanium™ Processor Event Qualification Modes (Continued)**

Event Qualification Modes	Instr. Address Range Check PMC[13].ta	Opcode Matching PMC[8]	Data Address Range Check PMC[11].pt
Instruction and Data Address Range Check	0	0xffff_ffff_ffff_ffff	0
Opcode Matching and Data Address Range Check	1	Desired Opcodes	0

### 6.1.3.2 Privilege Level Constraints

Performance monitoring software cannot always count on context switch support from the operating system. In general, this has made performance analysis of a single process in a multi-processing system or a multi-process workload very difficult. To provide hardware support for this kind of analysis, the Itanium architecture specifies three global bits (PSR.up, PSR.pp, DCR.pp) and a per-monitor “privilege monitor” bit (PMC[i].pm). To break down the performance contributions of operating system and user-level application components, each monitor specifies a 4-bit privilege level mask (PMC[i].plm). The mask is compared to the current privilege level in the processor status register (PSR.cpl), and event counting is enabled if PMC[i].plm[PSR.cpl] is one. The Itanium processor performance monitors control is discussed in [Section 6.2.1, "Performance Monitor Control and Accessibility"](#).

PMC registers can be configured as user-level monitors (PMC[i].pm is zero) or system-level monitors (PMC[i].pm is one). A user-level monitor is enabled whenever PSR.up is one. PSR.up can be controlled by an application using the `sum/rum` instructions. This allows applications to enable/disable performance monitoring for specific code sections. A system-level monitor is enabled whenever PSR.pp is one. PSR.pp can be controlled at privilege level 0 only, which allows monitor control without interference from user-level processes. The `pp` field in the default control register (DCR.pp) is copied into PSR.pp whenever an interruption is delivered. This allows events generated during interruptions to be broken down separately: if DCR.pp is zero, events during interruptions are not counted, if DCR.pp is one, they are included in the kernel counts.

As shown in [Figure 6-6](#), [Figure 6-7](#) and [Figure 6-8](#), single process, multi-process, and system level performance monitoring are possible by specifying the appropriate combination of PSR and DCR bits. These bits allow performance monitoring to be controlled entirely from a kernel level device driver, without explicit operating system support. Once the desired monitoring configuration has been setup in a process’ processor status register (PSR), “regular” unmodified operating context switch code automatically enables/disables performance monitoring.

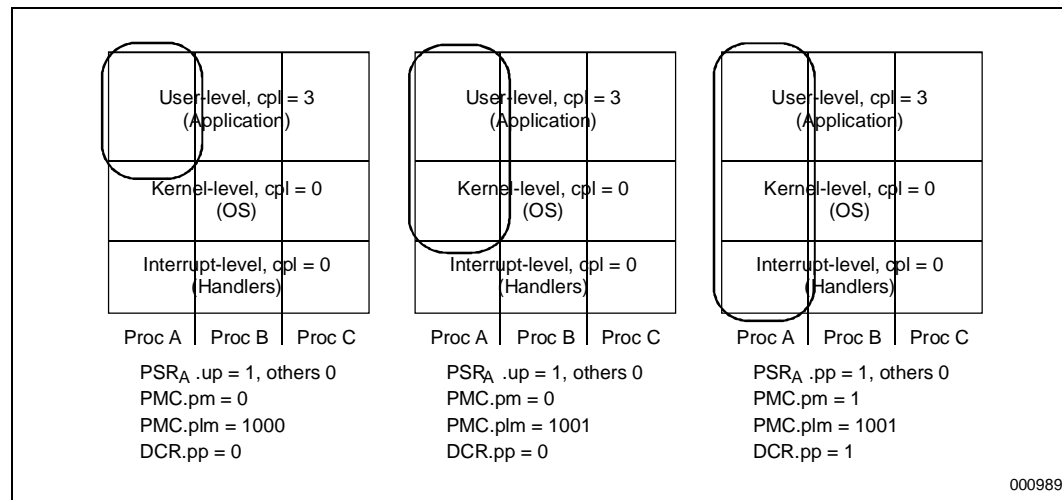
With support from the operating system, individual per-process break-down of event counts can be generated as outlined in [Section 7.2, "Performance Monitoring"](#) of [Volume 2 of the Intel® Itanium™ Architecture Software Developer’s Manual](#).

### 6.1.3.3 Instruction Set Constraints

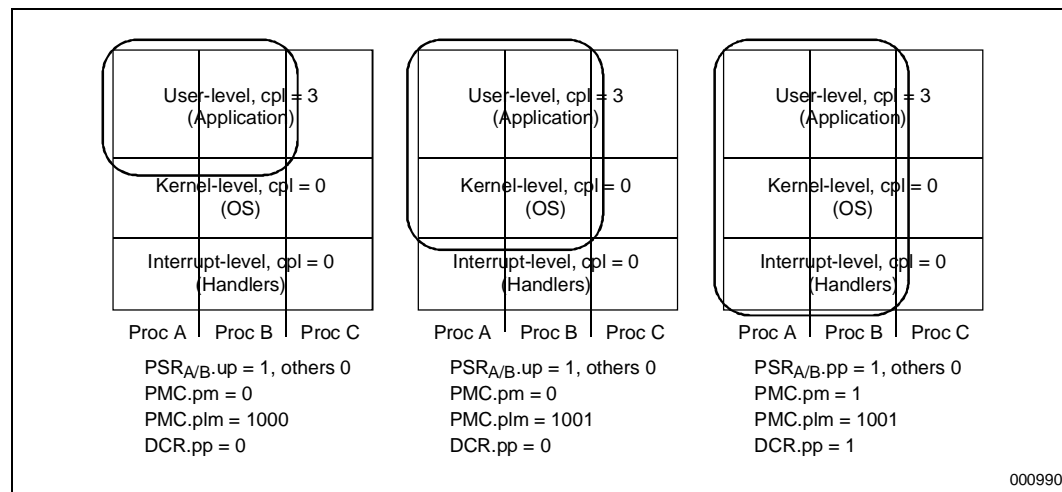
On the Itanium processor, monitoring can additionally be constrained based on the currently executing instruction set as defined by PSR.is. This capability is supported by the four generic performance counters as well as the instruction and data event address registers. However, the Itanium instruction address range checking, Itanium opcode matching and the Itanium branch trace buffer, only support Itanium-based code execution. When these Itanium architecture only features

are used, the corresponding PMC register instruction set mask (PMC[i].ism) should be set to Itanium architecture only (01) to ensure that events generated by IA-32 code do not corrupt the Itanium-based event counts.

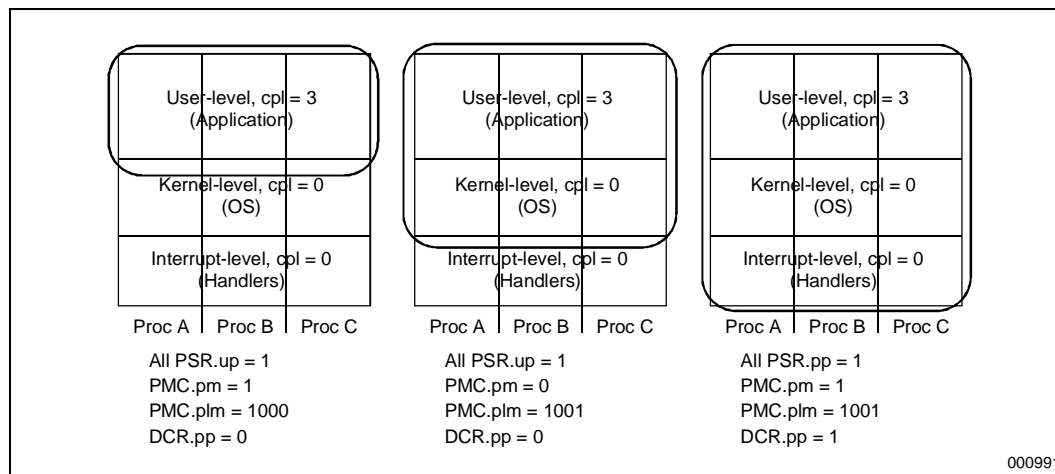
**Figure 6-6. Single Process Monitor**



**Figure 6-7. Multiple Process Monitor**



**Figure 6-8. System Wide Monitor**



## 6.2 Performance Monitor State

Two sets of performance monitor registers are defined. Performance Monitor Configuration (PMC) registers are used to configure the monitors. Performance Monitor Data (PMD) registers provide data values from the monitors. This section describes the Itanium processor performance monitoring registers which expands on the Itanium architectural definition. As shown in [Figure 6-9](#), the Itanium processor provides four 32-bit performance counters (PMC/PMD[4,5,6,7] pairs), and the following model-specific monitoring registers: instruction and data event address registers (EARs) for monitoring cache and TLB misses, a branch trace buffer, two opcode match registers and an instruction address range check register.

[Table 6-4](#) defines the PMC/PMD register assignments for each monitoring feature. The interrupt status registers are mapped to PMC[0,1,2,3]. The four generic performance counter pairs are assigned to PMC/PMD[4,5,6,7]. The event address registers and the branch trace buffer are controlled by three configuration registers (PMC[10,11,12]). Captured event addresses and cache miss latencies are accessible to software through five event address data registers (PMD[0,1,2,3,17]) and a branch trace buffer (PMD[8-16]). On the Itanium processor, monitoring of some events can additionally be constrained to a programmable instruction address range by appropriate setting of the instruction breakpoint registers (IBR) and the instruction address range check register (PMC[13]). Two opcode match registers (PMC[8,9]) allow monitoring of some events to be qualified with a programmable opcode. For memory operations, events can be qualified by a programmable data address range by appropriate setting of the data breakpoint registers (DBR) and the data address range check bits in PMC[11].

### 6.2.1 Performance Monitor Control and Accessibility

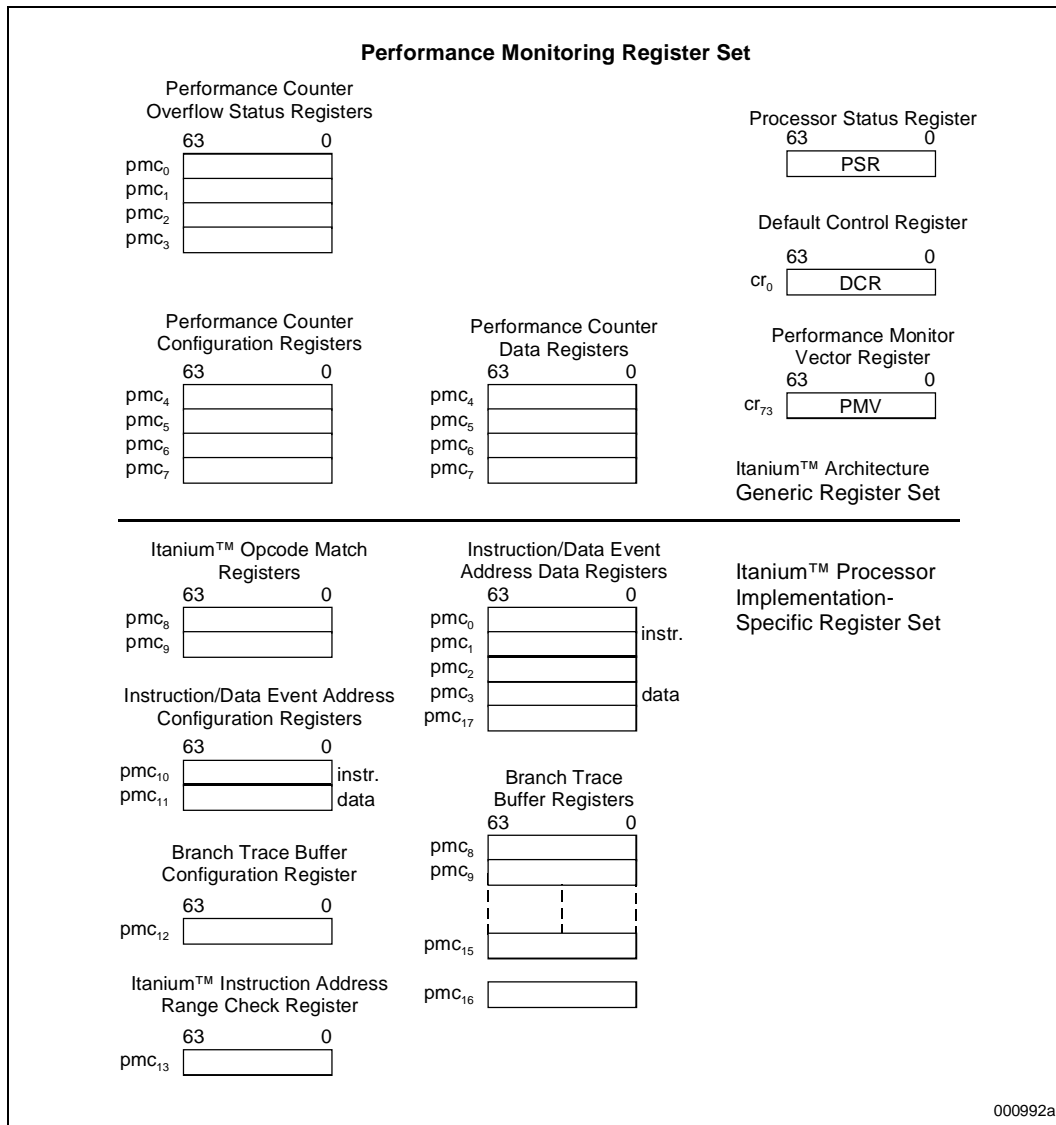
Event collection is controlled by the Performance Monitor Configuration (PMC) registers and the processor status register (PSR). Four PSR fields (PSR.up, PSR.pp, PSR.cpl and PSR.sp) and the performance monitor freeze bit (PMC[0].fr) affect the behavior of all performance monitor registers.

Finer, per monitor, control is provided by three PMC register fields (PMC[i].plm, PMC[i].ism, and PMC[i].pm). Instruction set masking based on PMC[i].ism is an Itanium processor model-specific feature. Event collection for a monitor is enabled under the following constraints on the Itanium processor:

$$\text{Monitor Enable}_i = (\text{not } \text{PMC}[0].\text{fr}) \text{ and } \text{PMC}[i].\text{plm}[\text{PSR}.\text{cpl}] \text{ and } ((\text{not } \text{PMC}[i].\text{ism}[\text{PSR}.\text{is}]) \text{ or } (\text{PMC}[i]=12)) \text{ and } (\text{not } (\text{PMC}[i].\text{pm}) \text{ and } \text{PSR}.\text{up}) \text{ or } (\text{PMC}[i].\text{pm} \text{ and } \text{PSR}.\text{pp})$$

Figure 3-2, “Processor Status Register (PSR)” on page 2:18 in Volume 2 of the *Intel® Itanium™ Architecture Software Developer’s Manual* defines the PSR control fields that affect performance monitoring. For a detailed definition of how the PSR bits affect event monitoring and control accessibility of PMD registers, please refer to Section 3.3.2, “Processor Status Register (PSR)” and Section 7.2.1, “Generic Performance Counter Registers” in Volume 2 of the *Intel® Itanium™ Architecture Software Developer’s Manual*.

**Figure 6-9. Intel® Itanium™ Processor Performance Monitor Register Model**




**Table 6-4. Intel® Itanium™ Processor Performance Monitor Register Set**

Monitoring Feature	Configuration Registers (PMC)	Data Registers (PMD)	Description
Interrupt Status	PMC[0,1,2,3]	none	See Section 6.2.3, "Performance Monitor Overflow Status Registers (PMC[0,1,2,3])"
Event Counters	PMC[4,5,6,7]	PMD[4,5,6,7]	See Section 6.2.2, "Performance Counter Registers"
Opcode Matching	PMC[8,9]	none	See Section 6.2.5, "Intel® Itanium™ Opcode Match Registers (PMC[8,9])"
Instruction EAR	PMC[10]	PMD[0,1]	See Section 6.2.7.1, "Instruction EAR (PMC[10], PMD[0,1])"
Data EAR	PMC[11]	PMD[2,3,17]	See Section 6.2.7.4, "Data EAR (PMC[11], PMD[2,3,17])"
Instruction Address Range Check	PMC[13]	none	See Section 6.2.4, "Intel® Itanium™ Instruction Address Range Check Register (PMC[13])"
Data Address Range Check	PMC[11]	none	See Section 6.2.6, "Intel® Itanium™ Data Address Range Check (PMC[11])"

As defined in Table 6-4, each of these PMC registers controls the behavior of its associated performance monitor data registers (PMD). Table 6-11 defines per monitor controls that apply to PMC[4,5,6,7,10,11,12]. The Itanium processor model-specific PMD registers associated with instruction/data EARs and the branch trace buffer (PMD[0,1,2,3,8-17]) can be read reliably only when event monitoring is frozen (PMC[0].fr is one).

**Figure 6-10. Processor Status Register (PSR) Fields for Performance Monitoring**



**Figure 6-11. Performance Monitor PMC Register Control Fields (PMC[4,5,6,7,10,11,12])**

Field	Bits	Description
plm	3:0	Privilege Level Mask - controls performance monitor operation for a specific privilege level. Each bit corresponds to one of the 4 privilege levels, with bit 0 corresponding to privilege level 0, bit 1 with privilege level 1, etc. A bit value of 1 indicates that the monitor is enabled at that privilege level. Writing zeros to all plm bits effectively disables the monitor. In this state, the Intel® Itanium™ processor will not preserve the value of the corresponding PMD register(s).
pm	6	Privileged monitor - When 0, the performance monitor is configured as a user monitor, and enabled by PSR.up. When PMC.pm is 1, the performance monitor is configured as a privileged monitor, enabled by PSR.pp, and PMD can only be read by privileged software.
ism	25:24	Instruction Set Mask - controls performance monitor operation based on the current instruction set. The instruction set mask applies to PMC[4,5,6,7,10,11] but not to PMC[12]. 00: monitoring enabled during Intel® Itanium™ instruction execution and IA-32 instruction execution (regardless of PSR.is) 10: bit 24 low enables monitoring during Intel® Itanium™ instruction execution (when PSR.is is zero) 01: bit 25 low enables monitoring during IA-32 instruction execution (when PSR.is is one) 11: disables monitoring


## 6.2.2 Performance Counter Registers

The Itanium processor provides four generic performance counters (PMC/PMD[4,5,6,7] pairs). The implemented counter width on the Itanium processor is 32 bits. The Itanium processor counters are not symmetrical (i.e. not all event types can be monitored by all counters). Counters PMC/PMD[4,5] can track events whose maximum per-cycle event increment is 7. Counters PMC/PMD[6,7] can track events whose maximum per-cycle event increment is 3.

The Itanium processor extends the generic Itanium counter configuration register (PMC) layout by adding two fields for specifying a unit mask (umask) and a threshold field. These model-specific fields are described in Table 6-5. A counter overflow occurs when the counter wraps (i.e. a carry out from bit 31 is detected). Software can force an external interruption or external notification after N events, by preloading the monitor with a count value of  $2^{32} - N$ . When accessible, software can continuously read the performance counter registers PMD[4,5,6,7] without disabling event collection. The processor guarantees that software will see monotonically increasing counter values.

Figure 6-12 and Table 6-5 define the layout of the Itanium processor Performance Counter Data Registers (PMD[4,5,6,7]). Figure 6-13, Figure 6-14 and Table 6-5 define the layout of the Itanium processor Performance Counter Configuration Registers (PMC[4,5,6,7]).


**Figure 6-12. Intel® Itanium™ Processor Generic PMD Registers (PMD[4,5,6,7])**




**Table 6-5. Intel® Itanium™ Processor Generic PMD Register Fields**

Field	Bits	Description
sxt32	63:32	Writes are ignored, Reads return the value of bit 31, so count values appear as sign extended.
count	31:0	Event Count. The counter is defined to overflow when the count field wraps (carry out from bit 31).

**Figure 6-13. Intel® Itanium™ Processor Generic PMC Registers (PMC[4,5])**



**Figure 6-14. Intel® Itanium™ Processor Generic PMC Registers (PMC[6,7])**




**Table 6-6. Intel® Itanium™ Processor Generic PMC Register Fields (PMC[4,5,6,7])**

Field	Bits	Description
plm	3:0	Privilege Level Mask. See <a href="#">Table 6-11, "Performance Monitor PMC Register Control Fields (PMC[4,5,6,7,10,11,12])"</a> .
ev	4	External visibility - When 1, an external notification (BPM pin strobe) is provided whenever the counter wraps, i.e. a carry out from bit 31 is detected. External notification occurs regardless of the setting of the oi bit. On the Intel® Itanium™ processor, PMC[4] external notification strobos the BPM0 pin, PMC[5] external notification strobos the BPM1 pin, PMC[6] external notification strobos the BPM2 pin, and PMC[7] external notification strobos the BPM3 pin.
oi	5	Overflow interrupt - When 1, a Performance Monitor Interrupt is raised and the performance monitor freeze bit (PMC[0].fr) is set when the monitor overflows. When 0, no interrupt is raised and the performance monitor freeze bit (PMC[0].fr) remains unchanged. Overflow occurs when the counter wraps, i.e. a carry out from bit 31 is detected. Counter overflows generate only one interrupt.
pm	6	Privilege Monitor. See <a href="#">Table 6-11, "Performance Monitor PMC Register Control Fields (PMC[4,5,6,7,10,11,12])"</a> .
ig	7	ignored
es	14:8	Event select - selects the performance event to be monitored. Intel® Itanium™ processor event encodings are defined in <a href="#">Chapter 7, "Performance Monitor Events"</a> .
ig	15	ignored
umask	19:16	Unit Mask - event specific mask bits (see event definition for details)
threshold	22:20 21:20	Threshold -enables thresholding for "multi-occurrence" events. PMC[4,5] define 3 threshold bits 22:20, while PMC[6,7] define 2 threshold bits 21:20. When threshold is zero, the counter sums up all observed event values. When the threshold is non-zero, the counter increments by one in every cycle in which the observed event value exceeds the threshold.
ism	25:24	Instruction Set Mask. See <a href="#">Table 6-11, "Performance Monitor PMC Register Control Fields (PMC[4,5,6,7,10,11,12])"</a> .
ignored	63:24	Read zero, Writes ignored.

### 6.2.3 Performance Monitor Overflow Status Registers (PMC[0,1,2,3])

The Itanium processor supports four counters. As shown in [Figure 6-15](#) and [Table 6-7](#) only PMC[0]{7:4} bits are populated. All other overflow bits are ignored, i.e. they read as zero and ignore writes.

**Figure 6-15. Intel® Itanium™ Processor Performance Monitor Overflow Status Registers (PMC[0,1,2,3])**




**Table 6-7. Intel® Itanium™ Processor Performance Monitor Overflow Register Fields (PMC[0,1,2,3])**

Register	Field	Bits	Description
PMC[0]	fr	0	Performance Monitor “freeze” bit - when 1, event monitoring is disabled. When 0, event monitoring is enabled. This bit is set by hardware whenever a performance monitor overflow occurs and its corresponding overflow interrupt bit (PMC.oi) is set to one. SW is responsible for clearing it. When the PMC.oi bit is not set, then counter overflows do not set this bit.
PMC[0]	ignored	3:1	Read zero, Writes ignored.
PMC[0]	overflow	7:4	Event Counter Overflow - When bit n is one, indicate that the PMD <sub>n</sub> overflowed. This is a bit vector indicating which performance monitor overflowed. These overflow bits are set on their corresponding counters overflow regardless of the state of the PMC.oi bit. These bits are sticky and multiple bits may be set.
PMC[0]	ignored	63:8	Read zero, Writes ignored.
PMC [1,2,3]	ignored	63:0	Read zero, Writes ignored.

## 6.2.4 Intel® Itanium™ Instruction Address Range Check Register (PMC[13])

The Itanium processor allows event monitoring to be constrained to a range of instruction addresses. All four architectural breakpoint registers (IBRs) are used to specify the desired address range. The Itanium processor instruction address range check register PMC[13] specifies how the resulting address match is applied to the performance monitors.

**Figure 6-16. Intel® Itanium™ Processor Instruction Address Range Check Register (PMC[13])**



**Table 6-8. Intel® Itanium™ Processor Instruction Address Range Check Register Fields (PMC[13])**

Field	Bits	Description
ta	0	Tag All - when 1, all events are counted independent of instruction address and instruction set. The default value of this PMC[13].ta should be set to one upon reset.

Instruction address range checking is controlled by the “tag all” bit (PMC[13].ta). When PMC[13].ta is one, all instructions are tagged regardless of IBR settings. In this mode, events from both IA-32 and Itanium-based code execution contribute to the event count. When PMC[13].ta is zero, the instruction address range check based on the IBR settings is applied to all Itanium-based code fetches. In this mode, IA-32 instructions are never tagged, and, as a result, events generated by IA-32 code execution are ignored. Table 6-9 defines the behavior of the instruction address range checker for different combinations of PSR.is and PMC[13].ta.



**Table 6-9. Intel® Itanium™ Processor Instruction Address Range Check by Instruction Set**

	PSR.is	
PMC <sub>13</sub> .ta	0 (Intel® Itanium™ Architecture)	1 (IA-32)
0	Tag only Intel® Itanium™ instructions if they match IBR range	DO NOT tag any IA-32 operations.
1	Tag all Intel® Itanium™ instructions and IA-32 instructions. Ignore IBR range.	

The processor compares every Itanium-based instruction fetch address IP{63:0} with each of the four architectural instruction breakpoint registers. Regardless of the value of the instruction break-point fault enable (IBR x-bit), the following expression is evaluated for each of the Itanium processor's four IBRs:

$$IBRmatch_i = \text{match}(\text{IP}, \text{IBR}_{[2*i]}.addr, \text{IBR}_{[2*i+1]}.mask, \text{IBR}_{[2*i+1]}.plm)$$

On the Itanium processor, in which only 54 virtual and 44 physical address bits are implemented, this IBR match is defined as follows:

$$IBRmatch_i = (\text{IBR}_{[2*i+1]}.plm[\text{PSR}.cpl])$$

$$\text{and } (\text{AND}_{b=50..0}((\text{IBR}_{[2*i]}.addr\{b\} \text{ and } \text{IBR}_{[2*i+1]}.mask\{b\}) = (\text{IP}\{b\} \text{ and } \text{IBR}_{[2*i+1]}.mask\{b\})))$$

$$\text{and } (\text{AND}_{b=55..51}((\text{IBR}_{[2*i]}.addr\{b\} \text{ and } \text{IBR}_{[2*i+1]}.mask\{b\}) = (\text{IP}\{50\} \text{ and } \text{IBR}_{[2*i+1]}.mask\{b\})))$$

$$\text{and } (\text{AND}_{b=60..56}(\text{IBR}_{[2*i]}.addr\{b\} = \text{IP}\{50\}))$$

$$\text{and } (\text{AND}_{b=63..61}(\text{IBR}_{[2*i]}.addr\{b\} = \text{IP}\{b\}))$$

The resulting four matches are combined with the PSR.is bit, two instruction address range check register bits, the IBR x-bits, and PSR.db:

$$IBRRangeTag = (\text{PMC}[13].ta)$$

$$\text{or } ((\text{not } \text{PSR}.is)$$

$$\text{and } ((\text{IBR}match_0 \text{ or } \text{IBR}match_1 \text{ or } \text{IBR}match_2 \text{ or } \text{IBR}match_3)$$

$$\text{and } (\text{not } (\text{PSR}.db \text{ or } \text{IBR}_1.x \text{ or } \text{IBR}_3.x \text{ or } \text{IBR}_5.x \text{ or } \text{IBR}_7.x))))$$

The instruction range check tag (IBRRangeTag) considers the IBR address ranges only if PMC[13].ta is zero, PSR.is is zero, and if none of the IBR x-bits or PSR.db are set. Since the architectural break-point registers (IBRs) are used to specify the desired performance monitor address range, it is not possible to constrain monitoring when the IBRs are used in their architectural break-point capacity, i.e. when PSR.db or an IBR x-bit is set. In other words, it is not possible to use performance monitor address range checking when a debugger is running, unless the debugger and the performance monitor software carefully synchronize their use of the IBRs.

The instruction range check tag is computed early in the processor pipeline and therefore includes speculative, wrong-path as well as predicated off instructions. Furthermore, range check tags are not accurate in the instruction fetch and out-of-order parts of the pipeline (cache and bus units). Therefore, software must accept a level of range check inaccuracy for events generated by these units, especially for non-looping code sequences that are shorter than the Itanium processor pipeline. As described in [Section 6.1.3.1, "Combining Opcode Matching, Instruction, and Data Address Range Check"](#), the instruction range check result may be combined with the results of the Itanium opcode match registers described in the next section.



As shown in [Figure 6-5](#), the two tags, Tag(PMC[8]) and Tag(PMC[9]), are staged down the processor pipeline until instruction retirement, and can be selected as a retired instruction count event. In this way, a performance counters (PMC/PMD[4,5,6,7]) can be used to count the number of retired instructions within the programmed range that match the specified opcodes. All combinations of the mifb bits are supported. To match A-syllable instructions both m and i bits should be set to one. To match all instruction types, all mifb and all mask bits should be set to one. This will count the number of retired instructions within the programmed address range. One of the combined Itanium address range and opcode match tags, Tag(PMC[8]), qualifies most down-stream pipeline events. To ensure that all events are counted independent of the Itanium opcode matcher, all mifb and all mask bits of PMC[8] should be set to one (all opcodes match). Tag(PMC[9]) is not used to qualify downstream events.

## 6.2.6 Intel® Itanium™ Data Address Range Check (PMC[11])

For instructions that reference memory, the Itanium processor allows event counting to be constrained by data address ranges using the architectural data breakpoint registers (DBRs). Data address range checking capability is controlled enabled by the “pass tags” bit in the Data Event Address Register (PMC[11].pt). For details on PMC[11], refer to [Section 6.2.7.4, "Data EAR \(PMC\[11\], PMD\[2,3,17\]\)"](#).

When enabled (PMC[11].pt is zero), data address range checking is applied to loads (all types), stores, semaphore operations, and the `lfetch` instruction whose upstream opcode match Tag(PMC[8]) was set. When PMC[11].pt is one, RSE operations and VHPT walks are tagged only if the opcode match Tag(PMC[8]) was set for the operation that caused the RSE or VHPT activity. When PMC[11].pt is zero, all RSE operations and VHPT walks that hit the programmed data address range are tagged (regardless of the opcode match Tag(PMC[8])). To capture all VHPT walks when PMC[11].pt is zero, the minimum DBR mask granularity must be set to the size of a single VHPT entry.

On the Itanium processor, in which only 54 virtual address bits are implemented, the performance monitoring DBR match function is defined as follows:

$DBRRangeMatch_i =$

$(AND_{b=50..0}((DBR_{[2*i]} .addr\{b\} \text{ and } DBR_{[2*i]+1} .mask\{b\}) = (addr\{b\} \text{ and } DBR_{[2*i]+1} .mask\{b\})))$

and  $(AND_{b=55..51}((DBR_{[2*i]} .addr\{b\} \text{ and } DBR_{[2*i]+1} .mask\{b\}) = (addr\{50\} \text{ and } DBR_{[2*i]+1} .mask\{b\})))$

$DBR_{[2*i]+1} .mask\{b\})))$

and  $(AND_{b=60..56}(DBR_{[2*i]} .addr\{b\} = addr\{50\}))$

and  $(AND_{b=63..61}(DBR_{[2*i]} .addr\{b\} = addr\{b\}))$

The resulting four matches are combined with PSR.db to form a single DBR match:

$DBRRangeMatch = ((DBRRangeMatch_0 \text{ or } DBRRangeMatch_1 \text{ or } DBRRangeMatch_2 \text{ or } DBRRangeMatch_3) \text{ and } (not \text{ PSR.db}))$

**Note:** DBR matching for performance monitoring ignores the setting of the DBR r, w and plm fields. Finally, the DBRRangeMatch is combined with PMC[11].pt and the upstream opcode match tag Tag(PMC[8]) as follows:

$$\text{DBRRangeTag} = \text{Tag}(\text{PMC}[8]) \text{ and } ((\text{PMC}[11].\text{pt}) \text{ or } \text{DBRRangeMatch})$$

DBR based data address range checking combined with opcode matching and instruction range checking allows the following combinations of event monitoring on the Itanium processor.


## 6.2.7 Event Address Registers (PMC[10,11]/PMD[0,1,2,3,17])

This section defines the register layout for the Itanium processor instruction and data event address registers (EARs). Sampling of four events is supported on the Itanium processor: instruction cache and instruction TLB misses, data cache load misses, and data TLB misses. The EARs are configured through two PMC registers (PMC[10,11]). EAR specific unit masks allow software to specify event collection parameters to hardware. Instruction and data addresses, operation latencies and other captured event parameters are provided in five PMD registers (PMD[0,1,2,3,17]). The instruction and data cache EARs report the latency of captured cache events and allow latency thresholding to qualify event capture. Event address data registers (PMD[0,1,2,3,17]) contain valid data only when event collection is frozen (PMC[0].fr is one). Reads of PMD[0,1,2,3,17] while event collection is enabled return undefined values.

### 6.2.7.1 Instruction EAR (PMC[10], PMD[0,1])

The instruction event address configuration register (PMC[10]) can be programmed to monitor either L1 instruction cache or instruction TLB miss events. [Figure 6-18](#) and [Table 6-11](#) detail the register layout of PMC[10]. [Figure 6-19](#) describes the associated event address data registers PMD[0,1].

**Figure 6-18. Instruction Event Address Configuration Register (PMC[10])**




When the tlb-bit (PMC[10].tlb) is set to zero instruction cache misses are monitored, when it is set to one instruction TLB misses are monitored. The interpretation of the umask field and performance monitor data registers PMD[0,1] depend on the setting of the tlb bit, and are described in [Section 6.2.7.2, "Instruction EAR Cache Mode \(PMC\[10\].tlb=0\)"](#) for instruction cache monitoring and in [Section 6.2.7.3, "Instruction EAR TLB Mode \(PMC\[10\].tlb=1\)"](#) for instruction TLB monitoring.

**Table 6-11. Instruction Event Address Configuration Register Fields (PMC[10])**

Field	Bits	Description
plm	3:0	See Table 6-11.
pm	6	See Table 6-11.
tlb	7	Instruction EAR selector: instruction cache/TLB
		if tlb=0: monitor L1 instruction cache misses PMD[0,1] register interpretation see Table 6-13.
		if tlb=1: monitor instruction TLB misses PMD[0,1] register interpretation see Table 6-15.
umask	19:16	Instruction EAR unit mask
		if tlb=0: instruction cache unit mask (definition see Table 6-12)
		if tlb=1: instruction TLB unit mask (definition see Table 6-14)
ism	25:24	See Table 6-11.

**Figure 6-19. Instruction Event Address Register Format (PMD[0,1])**



### 6.2.7.2 Instruction EAR Cache Mode (PMC[10].tlb=0)

When PMC[10].tlb is zero, the instruction event address register captures instruction addresses and access latencies for L1 instruction cache misses. Only misses whose latency exceeds a programmable threshold are captured. The threshold is specified as a four bit umask field in the configuration register PMC[10]. Possible threshold values are defined in Table 6-12.

As defined in Table 6-13, the address of the instruction cache line missed the L1 instruction cache is provided in PMD[0]. If no qualified event was captured, the valid bit in PMD[0] is zero. The latency of the captured instruction cache miss in processor clock cycles is provided in the latency field of PMD[1]. In cache mode, the TLB miss bit of PMD[0] is undefined.

**Table 6-12. Instruction EAR (PMC[10]) umask Field in Cache Mode (PMC[10].tlb=0)**

umask Bits 3:0	Latency Threshold [CPU cycles]	umask Bits 3:0	Latency Threshold [CPU cycles]
0000	>= 4	0110	>= 256
0001	>= 8	0111	>= 512
0010	>= 16	1000	>= 1024
0011	>= 32	1001	>= 2048
0100	>= 64	1010	>= 4096
0101	>= 128	1011.. 1111	No events are captured.

**Table 6-13. Instruction EAR (PMD[0,1]) in Cache Mode (PMC[10].tlb=0)**

Register	Field	Bits	Description
PMD[0]	v	0	Valid Bit 0: invalid address (EAR did not capture qualified event) 1: EAR contains valid event data
	tlb	1	TLB Miss Bit (undefined in cache mode)
	Instruction Cache Line Address	63:5	Address of instruction cache line that caused cache miss <sup>a</sup>
PMD[1]	latency	11:0	Latency in processor clocks

a. The Itanium processor does not implement virtual address bits va{60:51} and physical address bits pa{62:44}. The instruction and data address bits {60:51} of PMD[0] read as a sign-extension of bit {50}. Writes to bits {60:51} of PMD[0] are ignored by the processor.

### 6.2.7.3 Instruction EAR TLB Mode (PMC[10].tlb=1)

When PMC[10].tlb is one, the instruction event address register captures addresses of instruction TLB misses. The unit mask allows event address collection to capture specific subsets of instruction TLB misses. Table 6-14 summarizes the instruction TLB umask settings. All combinations of the mask bits are supported.

As defined in Table 6-15, the address of the instruction cache line fetch that missed the L1 TLB is provided in PMD[0]. The tlb bit indicates whether the captured TLB miss hit in the VHPT or required servicing by software. If no qualified event was captured, the valid bit in PMD[0] reads zero. In TLB mode, the latency field of PMD[1] is undefined.

**Table 6-14. Instruction EAR (PMC[10]) umask Field in TLB Mode (PMC[10].tlb=1)**

umask Bit	Instruction TLB EAR Unit Mask (Instruction TLB misses)
0	ignored
1	ignored
2	if one, capture Instruction TLB misses that hit VHPT
3	if one, capture Instruction TLB misses handled by software

**Table 6-15. Instruction EAR (PMD[0,1]) in TLB Mode (PMC[10].tlb=1)**


Register	Field	Bits	Description
PMD[0]	v	0	Valid Bit 0: invalid address (EAR did not capture qualified event) 1: EAR contains valid event data
	tlb	1	TLB Miss Bit: 0: VHPT Hit 1: Instruction TLB Miss handled by software
	Instruction Cache Line Address	63:5	Address of instruction cache line that caused TLB miss <sup>a</sup>
PMD[1]	latency	11:2	undefined in TLB mode

a. The Itanium processor does not implement virtual address bits va{60:51}. The instruction address bits {60:51} of PMD[0] read as a sign-extension of bit {50}. Writes to bits {60:51} of PMD[0] are ignored by the processor.

### 6.2.7.4 Data EAR (PMC[11], PMD[2,3,17])

The data event address configuration register (PMC[11]) can be programmed to monitor either L1 data cache load misses or L1 data TLB misses. Figure 6-20 and Table 6-16 detail the register layout of PMC[11]. Figure 6-21 describes the associated event address data registers PMD[2,3,17]. The tlb bit in configuration register PMC[11] selects data cache or data TLB monitoring. The interpretation of the umask field and registers PMD[2,3,17] depends on the setting of the tlb bit, and is described in Section 6.2.7.5, "Data Cache Load Miss Monitoring (PMC[11].tlb=0)" for data cache load miss monitoring and in Section 6.2.7.6, "Data TLB Miss Monitoring (PMC[11].tlb=1)" for data TLB monitoring. The PMC[11].pt bit controls data address range checking which is described in Section 6.2.6, "Intel® Itanium™ Data Address Range Check (PMC[11])".


**Figure 6-20. Data Event Address Configuration Register (PMC[11])**



**Table 6-16. Data Event Address Configuration Register Fields (PMC[11])**

Field	Bits	Description
plm	3:0	See Table 6-11.
pm	6	See Table 6-11.
tlb	7	Data EAR selector: data cache/TLB
		if tlb=0: monitor L1 data cache load misses PMD[2,3,17] register interpretation see Table 6-18.
		if tlb=1: monitor L1 data TLB misses PMD[2,3,17] register interpretation see Table 6-20.
umask	19:16	Data EAR unit mask if tlb=0: data cache unit mask (definition see Table 6-17) if tlb=1: data TLB unit mask (definition see Table 6-19)
ism	25:24	See Table 6-11.
pt	28	Pass Tags. This bit enables/disables data address range checking. See Section 6.2.6, "Intel® Itanium™ Data Address Range Check (PMC[11])" for details. if pt=1: then the Tag(PMC[8]) is passed down the pipeline unmodified. if pt=0: data address range checking is enabled for memory operations.

**Figure 6-21. Data Event Address Register Format (PMD[2,3,17])**



## 6.2.7.5 Data Cache Load Miss Monitoring (PMC[11].tlb=0)

If the Data EAR is configured to monitor data cache load misses (PMC[11].tlb=0), the umask is used as a load latency threshold defined by Table 6-17.

As defined in Table 6-18, the instruction and data addresses as well as the load latency of a captured data cache load miss is presented to software in three registers PMD[2,3,17]. If no qualified event was captured, the valid bit in PMD[3] is zero. In data cache load miss mode, the level field of PMD[3] is undefined.

**Table 6-17. PMC[11] Mask Fields in Data Cache Load Miss Mode (PMC[11].tlb=0)**

umask Bits 3:0	Latency Threshold [CPU cycles]	umask Bits 3:0	Latency Threshold [CPU cycles]
0000	>= 4	0110	>= 256
0001	>= 8	0111	>= 512
0010	>= 16	1000	>= 1024
0011	>= 32	1001	>= 2048
0100	>= 64	1010	>= 4096
0101	>= 128	1011.. 1111	No events are captured.

**Table 6-18. PMD[2,3,17] Fields in Data Cache Load Miss Mode (PMC[11].tlb=0)**

Register	Fields	Bit Range	Description
PMD[2]	Data Address	63:0	64-bit address of data item that caused miss <sup>a</sup>
PMD[3]	latency	11:0	Latency in CPU clocks
	level	63:62	Undefined in data cache load miss mode
PMD[17]	valid	0	Valid bit 0: invalid address (EAR did not capture qualified event) 1: EAR contains valid event data
	slot	3:2	Instruction bundle slot of memory instruction. For IA-32 ISA mode, this field is undefined.
	Instruction Address	63:4	Address of bundle that contains memory instruction. <sup>a</sup>

a. The Itanium processor does not implement virtual address bits va{60:51} and physical address bits pa{62:44}. The data/instruction address bits {60:51} of PMD[2,17] read as a sign-extension of bit {50}. Writes to bits {60:51} of PMD[2,17] are ignored by the processor.

The detection of data cache load misses requires a load instruction to be tracked during multiple clock cycles from instruction issue to cache miss occurrence. Since multiple loads may be outstanding at any point in time and the Itanium processor data cache miss event address register can only track a single load at a time, not all data cache load misses may be captured. When the processor hardware captures the address of a load (called the monitored load), it ignores all other overlapped concurrent loads until it is determined whether the monitored load turns out to be an L1 data cache miss or not. If the monitored load turns out to be a cache miss, its parameters are latched into PMD[2,3,17]. The processor randomizes the choice of which load instructions are tracked to prevent the same data cache load miss from always being captured (in a regular sequence of overlapped data cache load misses). While this mechanism will not always capture all data cache load misses in a particular sequence of overlapped loads, its accuracy is sufficient to be used by statistical sampling or code instrumentation.



### 6.2.7.6 Data TLB Miss Monitoring (PMC[11].tlb=1)

If the Data EAR is configured to monitor data TLB misses (PMC[11].tlb=1), the umask defined by Table 6-19 determine which data TLB misses are captured by the Data EAR. For TLB monitoring, all combinations of the mask bits are supported.

**Table 6-19. PMC[11] Unmask Field in TLB Miss Mode (PMC[11].tlb=1)**

umask Bit	Data EAR Unit Mask (L1 data TLB misses)
0	reserved
1	if one, capture L1 TLB misses that hit L2 Data TLB
2	if one, capture L1 TLB misses that hit VHPT
3	if one, capture L1 TLB misses that was handled by software

As defined in Table 6-20, the instruction and data addresses of captured data TLB misses are presented to software in PMD[2,17]. The level of the TLB hierarchy from which the L1 data TLB miss was satisfied is recorded in the level field of PMD[3]. If no qualified event was captured, the valid bit in PMD[17] and the level field in PMD[3] read zero. When programmed for data TLB monitoring, the contents of the latency field of PMD[3] are undefined.

**Table 6-20. PMD[2,3,17] Fields in TLB Miss Mode (PMC[11].tlb=1)**

Register	Field	Bit Range	Description
PMD[2]	Data Address	63:0	64-bit address of data item that caused miss <sup>a</sup>
PMD[3]	latency	11:0	Undefined in TLB Miss mode
	level	63:62	Data TLB Miss Level 0: invalid address (EAR did not capture qualified event) 1: L2 Data TLB hit 2: VHPT hit 3: Data TLB miss handled by software
PMD[17]	valid	0	Valid Bit: 0: invalid address (EAR did not capture qualified event) 1: EAR contains valid event data
	slot	3:2	Instruction Bundle Slot of memory instruction. In IA-32 ISA mode, this field is undefined.
	Instruction Address	63:4	Address of bundle that contains memory instruction <sup>a</sup>

a. The Itanium processor does not implement virtual address bits va{60:51} and physical address bits pa{62:44}. The data/instruction address bits {60:51} of PMD[2,17] read as a sign-extension of bit {50}. Writes to bits {60:51} of PMD[2,17] are ignored by the processor.

### 6.2.8 Intel® Itanium™ Branch Trace Buffer

The branch trace buffer provides information about the outcome of the most recent Itanium branch instructions and their predictions and outcomes. The Itanium branch trace buffer configuration register (PMC[12]) defines the conditions under which branch instructions are captured and allows the trace buffer to capture specific subsets of branch events. The Itanium branch trace buffer operates only during Itanium-based code execution (i.e. when PSR.is is zero).

In every cycle in which a qualified Itanium branch retires, its source bundle address and slot number are written to the branch trace buffer. The branches' target address is written to the next buffer location. If the target instruction bundle itself contains a qualified Itanium branch, the branch


trace buffer either records a single trace buffer entry (with the b-bit set) or makes two trace buffer entries: one that records the target instruction as a branch target (b-bit cleared), and another that records the target instruction as a branch source (b-bit set). As a result, the branch trace buffer may contain a mixed sequence of the branches and targets.

### 6.2.8.1 Intel® Itanium™ Trace Buffer Collection Constraining

The Itanium branch trace buffer configuration register (PMC[12]) defines the conditions under which branch instructions are captured. These conditions are given in Figure 6-22 and Table 6-21, and refer to conditions associated with the branch prediction and resolution hardware. These conditions are:

- Which branch prediction hardware structure made the prediction,
- The path of the branch (not taken/taken),
- Whether or not the branch path was mispredicted, and
- Whether or not the target of the branch was mispredicted.

**Figure 6-22. Intel® Itanium™ Branch Trace Buffer Configuration Register (PMC[12])**



**Table 6-21. Intel® Itanium™ Branch Trace Buffer Configuration Register Fields (PMC[12])**

Field	Bits	Description
plm	3:0	See Table 6-11.
pm	6	See Table 6-11.
tar	7	Target Address Register: 1: capture TAR predictions 0: No TAR predictions are captured
tm	9:8	Taken Mask: 11: all Intel® Itanium™ branches 10: Taken Intel® Itanium™ branches only 01: Not Taken Intel® Itanium™ branches only 00: No branch is captured
ptm	11:10	Predicted Target Address Mask: 11: capture branch regardless of target prediction outcome 10: branch predicted target address correctly 01: branch mispredicted target address 00: No branch is captured
ppm	13:12	Predicted Predicate Mask: 11: capture branch regardless of predicate prediction outcome 10: branch predicted branch path (taken/not taken) correctly 01: branch mispredicted branch path (taken/not taken) 00: No branch is captured
bpt	14	Branch Prediction Table: 10: No TAC predictions are captured
bac	15	Branch Address Calculator: 1: capture BAC predictions 0: No BAC predictions are captured



The Itanium processor uses the following micro-architectural structures for branch prediction: the Target Address Registers (TAR), and Target Address Cache (TAC). Using the tar and bac fields of the branch trace buffer configuration register (PMC[12]), collection in the branch trace buffer can be restricted to only branches predicted by a subset of these prediction structures.

The Target Address Registers (TAR) are a small and fast fully associative buffer that is exclusively written to by branch predict instructions with the '.imp' extension. A hit in the TAR will cause a taken prediction and yield the target address of the branch. If the tar field in the branch trace buffer configuration register (PMC[12]) is set to one, branches predicted by TAR will be included in the trace buffer.

The Target Address Cache (TAC) is a larger structure that is also written to by branch predict instructions, or the prediction hardware. The primary function of the TAC is to provide the target address of a branch.

If the bpt field in the branch trace buffer configuration register (PMC[12]) is set to one, branches predicted by the TAC will be included in the trace buffer.

If neither the TAR nor TAC generated a hit, the branch has to be predicted using the static hints encoded in the branches and the target address has to be calculated. This is done by the branch address corrector (BAC). If the bac field in the branch trace buffer configuration register (PMC[12]) is set to one, branches predicted by the branch address corrector will be included in the trace buffer.

Furthermore, using the ptm, ppm and tm fields in the branch trace buffer configuration register (PMC[12]) collection in the branch trace buffer can be restricted based on the correctness of target and predicate prediction in addition to whether the branch was actually taken or not.

To summarize, an Itanium branch and its target are captured by the trace buffer if the following equation is true:


```
(not PSR.is)
and (      (tm[1] and branch taken)
          or (tm[0] and branch not taken)
      )
and (      (ptm[1] and hardware predicted target address correctly
            and hardware predicted the branch path correctly
            and branch is taken)
          or (ptm[0] and hardware mispredicted target address
            and hardware predicted the branch path correctly
            and branch is taken)
          or (ptm[0] and ptm[1])
      )
and (      (ppm[1] and hardware predicted the branch path correctly)
          or (ppm[0] and hardware mispredicted the branch path)
      )
and (      (bpt and branch was predicted by TAC)
          or (bac and branch was predicted by BAC)
          or (tar and branch was predicted by TAR)
      )
)
```

To capture all mispredicted Itanium branches, the branch trace buffer configuration settings in PMC[12] should be: Tm=11, ptm=01, ppm=01, bpt=1, bac=1, and tar=1.

## 6.2.8.2 Intel® Itanium™ Branch Trace Buffer Reading

The eight branch trace buffer registers PMD[8-15] provide information about the outcome of a captured branch sequence. The branch trace buffer registers (PMD[8-15]) contain valid data only when event collection is frozen (PMC[0].fr is one). While event collection is enabled, reads of PMD[8-15] return undefined values. The registers follow the layout defined in Figure 6-23, and contain the address of either a captured branch instruction (b-bit=1) or branch target (b-bit=0). For branch instructions, the mp-bit indicates a branch misprediction. A branch trace register with a zero b-bit and a zero mp-bit indicates an invalid branch trace buffer entry. The slot field captures the slot number of the first taken Itanium branch instruction in the captured instruction bundle. A slot number of 3 indicates a not-taken branch. The target address bundle of a branch to IA-32 (br . ia) is recorded. An IA-32 JMPE branch instruction and its Itanium-based target are not recorded.

**Figure 6-23. Branch Trace Buffer Register Format (PMD[8-15])**



**Table 6-22. Intel® Itanium™ Branch Trace Buffer Register Fields (PMD[8-15])**

Field	Bit Range	Description
b	0	Branch Bit 1: contents of register is a branch instruction 0: contents of register is a branch target
mp	1	Mispredict Bit if b=1 and mp=1: mispredicted branch (due to target or predicate misprediction) if b=1 and mp=0: correctly predicted branch if b=0 and mp=0: <b>invalid branch trace buffer register</b> if b=0 and mp=1: valid target address
slot	3:2	if b=0: 00 if b=1: Slot index of first taken branch instruction in bundle 00: Itanium™-based Slot 0 branch/target 01: Itanium™-based Slot 1 branch/target 10: Itanium™-based Slot 2 branch/target 11: this was a not taken branch
Address	63:4	if b=1: 60-bit bundle address of Intel® Itanium™ branch instruction <sup>a</sup> if b=0: 60-bit target bundle address of Intel® Itanium™ branch instruction <sup>a</sup>

a. The Itanium processor does not implement virtual address bits va{60:51} and physical address bits pa{62:44}. When the processor captures an instruction address, bits {60:51} of PMD[8-15] are written by the processor with a sign-extension of bit {50} of the captured address. When PMD[8-15] are written by software bits {60:51} of PMD[8-15] can be written with any value (not necessarily a sign-extension of bit {50}).

In every cycle in which a qualified Itanium branch retires<sup>1</sup>, its source bundle address and slot number are written to the branch trace buffer. The branches' target address is written to the next buffer location. If the target instruction bundle itself contains a qualified Itanium branch, the branch trace buffer either records a single trace buffer entry (with the b-bit set) or makes two trace buffer

1. In some cases, the Itanium processor branch trace buffer will capture the source (but not the target) address of an excepting branch instruction. This occurs on trapping branch instructions as well as faulting br . ia, break . b and multiway branches.


entries: one that records the target instruction as a branch target (b-bit cleared), and another that records the target instruction as a branch source (b-bit set). As a result, the branch trace buffer may contain a mixed sequence of the branches and targets.

The Itanium branch trace buffer is a circular buffer containing the last four to eight qualified Itanium branches. The Branch Trace Buffer Index Register (PMD[16]) defined in Figure 6-24 identifies the most recently recorded branch or target. In every cycle in which a qualified branch (branch or target) is recorded, the branch buffer index (bbi) is post-incremented. After 8 entries have been recorded, the branch index wraps around, and the next qualified branch will overwrite the first trace buffer entry. The wrap condition itself is recorded in the full bit of PMD[16]. The bbi field of PMD[16] defines the next branch buffer index that is about to be written. The following formula computes the last written branch trace buffer PMD index from the contents of PMD[16]:

$$\text{last-written-PMD-index} = 8 + ((8 * \text{PMD}[16].\text{full}) + (\text{PMC}[16].\text{bbi} - 1) \% 8)$$

If both the full bit and the bbi field of PMD[16] are zero, no qualified branch has been captured by the branch trace buffer. The full bit gets set the every time the branch trace buffer wraps from PMD[15] to PMD[8]. Once set, the full bit remains set until explicitly cleared by software, i.e. it is a sticky bit. Software can reset the bbi index and the full bit by writing to PMD[16].

**Figure 6-24. Intel® Itanium™ Branch Trace Buffer Index Register Format (PMD[16])**



**Table 6-23. Intel® Itanium™ Branch Trace Buffer Index Register Fields (PMD[16])**

Field	Bit Range	Description
bbi	2:0	Branch Buffer Index [Range 0..7 - Index 0 indicates PMD[8]] Pointer to the next branch trace buffer entry to be written. if full=1: points to the oldest recorded branch/target if full=0: points to the next location to be written
full	3	Full Bit (sticky) if full=1: branch trace buffer has wrapped if full=0: branch trace buffer has not wrapped

## 6.2.9 Processor Reset, PAL Calls, and Low Power State

**Processor Reset:** On processor hardware reset bits oi, ev of all PMC registers are zero, and PMV.m is set to one. This ensures that no interrupts are generated, and events are not externally visible. On reset, PAL firmware ensures that the instruction address range check, the opcode matcher and the data address range check are initialized as follows:

- PMC[13].ta=1,
- PMC[8,9].mifb=1111, PMC[8,9].mask{29:3}= “all 1s”, PMC[8,9].match{59:33}= “all 0s”, and
- PMC[11].pt is 1.

All other performance monitoring related state is undefined.

**PAL Call:** As defined in Chapter 11, “Processor Abstraction Layer” in Volume 2 of the *Intel® Itanium™ Architecture Software Developer’s Manual*, the PAL call PAL\_PERF\_MON\_INFO provides software with information about the implemented performance monitors. The Itanium processor specific values are summarized in Table 6-24.

**Low Power State:** To ensure that monitor counts are preserved when the processor enters low power state, PAL\_LIGHT\_HALT freezes event monitoring prior to powering down the processor. PAL\_LIGHT\_HALT preserves the original value of the PMC[0] register.

**Table 6-24. Information Returned by PAL\_PERF\_MON\_INFO for the Intel® Itanium™ Processor**

PAL_PERF_MON_INFO Return Value	Description	Intel® Itanium™ Processor- specific Value
PAL_RETIRED	8-bit unsigned event type for counting the number of untagged retired Intel® Itanium™ instructions.	0x08
PAL_CYCLES	8-bit unsigned event type for counting the number of running CPU cycles.	0x12
PAL_WIDTH	8-bit unsigned number of implemented counter bits.	32
PAL_GENERIC_PM_PAIRS	8-bit unsigned number of generic PMC/PMD pairs.	4
PAL_PMCmask	256-bit mask defining which PMC registers are populated.	0x3FFF
PAL_PMDmask	256-bit mask defining which PMD registers are populated.	0x3FFFF
PAL_CYCLES_MASK	256-bit mask defining which PMC/PMD counters can count running CPU cycles (event defined by PAL_CYCLES)	0xF0
PAL_RETIRED_MASK	256-bit mask defining which PMC/PMD counters can count untagged retired Intel® Itanium™ instructions (event defined by PAL_RETIRED)	0x10

## 6.2.10 References

- [gprof] S.L. Graham S.L., P.B. Kessler and M.K. McKusick, “gprof: A Call Graph Execution Profiler”, Proceedings SIGPLAN’82 Symposium on Compiler Construction; SIGPLAN Notices; Vol. 17, No. 6, pp. 120-126, June 1982.
- [Lebeck] Alvin R. Lebeck and David A. Wood, “Cache Profiling and the SPEC benchmarks: A Case Study”, Tech Report 1164, Computer Science Dept., University of Wisconsin - Madison, July 1993.
- [VTune] Mark Atkins and Ramesh Subramaniam, “PC Software Performance Tuning”, IEEE Computer, Vol. 29, No. 8, pp. 47-54, August 1996.
- [WinNT] Russ Blake, “Optimizing Windows NT™”, Volume 4 of the Microsoft “Windows NT Resource Kit for Windows NT Version 3.51”, Microsoft Press, 1995.

This chapter describes the architectural and microarchitectural events on the Itanium processor whose occurrences are countable through the performance monitoring mechanisms described earlier in [Chapter 6](#). The earlier sections of this chapter aim to provide a high-level view of the event list, grouping logically related events together. Computation (either directly by a counter in hardware, or indirectly as a “derived” event) of common performance metrics is also discussed. Each directly measurable event is then described in greater detail in the alphabetized list of all processor events in [Section 7.8](#), “Performance Monitor Event List”.

## 7.1 Categorization of Events

Performance related events are grouped into the following categories:

- Basic Events: clock cycles, retired instructions ([Section 7.2](#))
- Instruction Execution: instruction decode, issue and execution, data and control speculation, and memory operations ([Section 7.3](#))
- Cycle Accounting Events: stall and execution cycle breakdowns ([Section 7.4](#))
- Branch Events: branch prediction ([Section 7.5](#))
- Memory Hierarchy: instruction and data caches ([Section 7.6](#))
- System Events: operating system monitors, instruction and data TLBs ([Section 7.7](#))

Each section listed above includes a table of all events (both directly measured and derived) in that category. Directly measurable events often use the PMC.umask field (See [Table 6-7](#) in [Chapter 6](#)) to measure a certain variant of the event in question. Symbolic event names for such events (e.g. ALAT\_REPLACEMENT.ALL) include a period to indicate use of the umask, specified by 4 bits in the detailed event description (x’s are for don’t-cares). Derived events are computable from directly measured events and include a “.d” suffix in their symbolic event names. Formulas to compute relevant derived events also appear in each section. Derived events are not, however, discussed in the systematic event listing in [Section 7.8](#).

The tables in the subsequent sections define events by specifying three attributes: symbolic event name, a brief event title and a reference to the detailed event description page. Derived events are not listed in the detailed event description pages and hence lack the appropriate reference.

## 7.2 Basic Events

[Table 7-1](#) summarizes four basic execution monitors. The CPU\_CYCLES event can be used to break out separate or combined Itanium or IA-32 cycle counts (by constraining the PMC/PMD based on the currently executing instruction set). The Itanium retired instruction count (IA64\_INST\_RETIRED) includes predicated true and false instructions, and nops, but excludes RSE operations.

**Table 7-1. Intel® Itanium™ and IA-32 Instruction Set Execution and Retirement Monitors**

Execution Monitors	Title
CPU_CYCLES	CPU Cycles on page 88.
IA64_INST_RETIRED	Retired Itanium Instructions on page 92.
IA32_INST_RETIRED	Retired IA-32 Instructions on page 92.
ISA_TRANSITIONS	Itanium ISA to IA-32 ISA Transitions on page 94.

Table 7-2 defines IPC and average instructions/cycles per ISA transition metrics.

**Table 7-2. Intel® Itanium™ and IA-32 Instruction Set Execution and Retirement Performance Metrics**

Performance Metric	Performance Monitor Equation
Intel® Itanium™ Instruction per Cycle	$IA64\_INST\_RETIRED / CPU\_CYCLES$ [Intel® Itanium™ only]
IA-32 Instruction per Cycle	$IA32\_INST\_RETIRED / CPU\_CYCLES$ [IA-32 only]
Average Intel® Itanium™ Instructions/Transition	$IA64\_INST\_RETIRED / (ISA\_TRANSITIONS * 2)$
Average IA-32 Instructions/Transition	$IA32\_INST\_RETIRED / (ISA\_TRANSITIONS * 2)$
Average Intel® Itanium™ Cycles/Transition	$CPU\_CYCLES[IA64] / (ISA\_TRANSITIONS * 2)$
Average IA-32 Cycles/Transition	$CPU\_CYCLES[IA32] / (ISA\_TRANSITIONS * 2)$

## 7.3 Instruction Execution

This section describes events related to instruction issue and retirement (Table 7-3, Table 7-4) multi-media and FP (Table 7-5), data and control speculation (Table 7-7), as well as memory monitors (Table 7-9).

**Table 7-3. Instruction Issue and Retirement Events**

Decode, Issue, Retirement Monitors	Description
INST_DISPERSED	Instructions Dispersed on page 93.
EXPL_STOPS_DISPERSED	Explicit Stops Dispersed on page 91.
ALL_STOPS_DISPERSED	Implicit and Explicit Stops Dispersed on page 67.
IA64_TAGGED_INST_RETIRED	Retired Tagged Itanium Instructions on page 92.
NOPS_RETIRED	Retired Nop Instructions on page 105.
PREDICATE_SQUASHED_RETIRED	Instructions Squashed Due to Predicate Off on page 106.
RSE_REFERENCES_RETIRED	RSE Accesses on page 107.
RSE_LOADS_RETIRED	RSE Load Accesses on page 106.

**Table 7-4. Instruction Issue and Retirement Events (Derived)**

Decode, Issue, Retirement Monitors	Description	Performance Monitor Equation
RSE_STORES_RETIRED.d	RSE Store Accesses	$RSE\_REFERENCES\_RETIRED - RSE\_LOADS\_RETIRED$

Instruction cache lines are delivered to the execution core and are dispersed to the Itanium processor functional units. The number of dispersed instructions (INST\_DISPERSED) on every cycle depends on the stops in the instruction stream (EXPL\_STOPS\_DISPERSED) as well as



functional unit availability. Resource limitations and branch bundles (regardless of prediction) force a break in the instruction dispersal. Therefore, they are known as implicit stops, and can be computed as `ALL_STOPS_DISPERSSED - EXPL_STOPS_DISPERSSED`.

Retired instruction counts (`IA64_TAGGED_INST_RETIRED`, `NOPS_RETIRED`) are based on tag information specified by the address range check and opcode match facilities. The tagged retired instruction counts include predicated off instructions. A separate event (`PREDICATE_SQUASHED_RETIRED`) is provided to count predicated off instructions. `RSE_REFERENCES_RETIRED` counts the number of retired RSE operations.

There are two ways to count the total number of retired Itanium instructions. Either the untagged `IA64_INST_RETIRED` event can be used or the `IA64_TAGGED_INST_RETIRED` event can be used by setting up the `PMC8` opcode match register to its don't care setting.

The FP monitors listed in [Table 7-5](#) (`FP_SIR_FLUSH`, `FP_FLUSH_TO_ZERO`) capture dynamic information about pipeline flushes and flush-to-zero occurrences due to floating-point operations. `FP_OPS_RETIRED.d` is a derived event that counts the number of retired FP operations.

**Table 7-5. Floating-point Execution Monitors**

Floating-point Monitors	Description
<code>FP_FLUSH_TO_ZERO</code>	FP Result Flushed to Zero on page 91.
<code>FP_SIR_FLUSH</code>	FP SIR Flushes on page 92.

**Table 7-6. Floating-point Execution Monitors (Derived)**

Floating-Point Monitors	Description	Performance Monitor Equation
<code>FP_OPS_RETIRED.d</code>	FP Operations Retired	$(4 * \text{FP\_OPS\_RETIRED\_HI}) + \text{FP\_OPS\_RETIRED\_LO}$

As [Table 7-7](#) describes, monitors for control and data speculation capture dynamic run-time information: the number of failed chk.s instructions (`INST_FAILED_CHKS_RETIRED.ALL`), the number of advanced load checks and check loads (`ALAT_INST_CHKA_LDC.ALL`) and failed advanced load checks and check loads (`ALAT_INST_FAILED_CHKA_LDC.ALL`) as seen by the ALAT. The number of retired chk.s instructions is monitored by the `IA64_TAGGED_INST_RETIRED` event with the appropriate opcode mask. Since the Itanium processor ALAT is updated by operations on mispredicted branch paths the number of advanced load checks and check loads needs an explicit event (`ALAT_INST_CHKA_LDC.ALL`). Finally, the `ALAT_REPLACEMENT.ALL` event can be used to monitor ALAT overflows.

**Table 7-7. Control and Data Speculation Monitors**

Control and Data Speculation Monitors	Description
<code>INST_FAILED_CHKS_RETIRED.ALL</code>	Failed Speculative Check Loads on page 93.
<code>ALAT_INST_CHKA_LDC.ALL</code>	Advanced Load Checks and Check Loads on page 65.
<code>ALAT_INST_FAILED_CHKA_LDC.ALL</code>	Failed Advanced Load Checks and Check Loads on page 66.
<code>ALAT_REPLACEMENT.ALL</code>	ALAT Entries Replaced by Any Instruction on page 64.

Using an instruction type unit mask the four control and data speculation events can be constrained to monitor integer, FP or all speculative instructions. With the Itanium processor speculation monitors, the performance metrics described in [Table 7-8](#) can be computed.

**Table 7-8. Control/Data Speculation Performance Metrics**

Performance Metric	Performance Monitor Equation
Control Speculation Miss Ratio	INST_FAILED_CHKS_RETIRED.ALL / (IA64_TAGGED_INST_RETIRED[chk.s]-PREDICATE_SQUASHED_RETIRED[chk.s])
Data Speculation Miss Ratio	ALAT_INST_FAILED_CHKA_LDC.ALL / ALAT_INST_CHKA_LDC.ALL
ALAT Capacity Miss Ratio	ALAT_REPLACEMENT.ALL / IA64_TAGGED_INST_RETIRED[id.a,ld.sa,ld.c.nc,ldf.a,ldf.sa,ldf.c.nc]-PREDICATE_SQUASHED_RETIRED[id.a,ld.sa,ld.c.nc,ldf.a,ldf.sa,ldf.c.nc])

The equations described in [Table 7-8](#) for Control Speculation Miss Ratio and ALAT Capacity Miss Ratio involve subtracting PREDICATE\_SQUASHED\_RETIRED[some inst] from IA64\_TAGGED\_INST\_RETIRED[some inst]. This is done because IA64\_TAGGED\_INST\_RETIRED includes predicated off instructions in its count, which do not update architectural state and hence need to be discounted in computing any performance metric. Using the opcode matcher in PMC8 with PREDICATE\_SQUASHED\_RETIRED (along with IA64\_TAGGED\_INST\_RETIRED) allows us to count the number of predicated off instances of that instruction as well. Note that computing the ALAT Capacity Miss Ratio will require multiple runs in order to obtain all the terms in the equation. This is done to the limitations imposed by the opcode matcher.

Finally, [Table 7-9](#) defines six memory instruction retirement events to count retired loads and stores. These counts include RSE operations. The load counts include failed check load instructions.

**Table 7-9. Memory Events**

Memory Monitors	Description
LOADS_RETIRED	Retired Loads on page 104.
STORES_RETIRED	Retired Stores on page 107.
UC_LOADS_RETIRED	Retired Uncacheable Loads on page 107.
UC_STORES_RETIRED	Retired Uncacheable Stores on page 107.
MISALIGNED_LOADS_RETIRED	Retired Unaligned Load Instructions on page 104.
MISALIGNED_STORES_RETIRED	Retired Unaligned Store Instructions on page 105.

## 7.4 Cycle Accounting Events

As described in [Section 6.1.1.4, “Cycle Accounting”](#), the Itanium processor provides eight directly measured stall cycle monitors. [Table 7-10](#) lists the cycle accounting events.

The Itanium processor classifies every clock cycle into one of 4 cycle counters, namely DEPENDENCY\_ALL\_CYCLE, MEMORY\_CYCLE, UNSTALLED\_BACKEND\_CYCLE, and PIPELINE\_ALL\_FLUSH\_CYCLE. The values of these 4 counters should add up to CPU\_CYCLES.

DEPENDENCY\_ALL\_CYCLE counts the number of cycles lost to instruction dispersal breaks (including both explicit and implicit stops), FP-related flushes and scoreboard stalls on GR or FR dependencies on non-load instructions. That is, the monitor does not count stalls that occur when an

instruction is waiting for source operands from the memory subsystem. Also note that this monitor does not count the number of cycles when the machine is executing instructions without stalls or flushes. The `DEPENDENCY_SCOREBOARD_CYCLE` monitor operates similarly, but does not include instruction dispersal breaks.

**Table 7-10. Stall Cycle Monitors**

<b>Stall Accounting Monitors</b>	<b>Description</b>
<code>PIPELINE_BACKEND_FLUSH_CYCLE</code>	Combination of Pipeline Flush Cycles caused by either a Branch Misprediction or an ExceptionCategory: Stall on page 105.
<code>DATA_ACCESS_CYCLE</code>	Data Access Stall Cycles on page 89.
<code>DEPENDENCY_SCOREBOARD_CYCLE</code>	Scoreboard Dependency Cycles on page 90.
<code>INST_ACCESS_CYCLE</code>	Instruction Access Cycles on page 93.
<code>PIPELINE_ALL_FLUSH_CYCLE</code>	Combination of Pipeline Flush Cycles caused by either a front-end or a back-end source on page 105.
<code>MEMORY_CYCLE</code>	Combined Memory Stall Cycles on page 104.
<code>DEPENDENCY_ALL_CYCLE</code>	Scoreboard Dependency and Dispersal Break Cycles on page 89.
<code>UNSTALLED_BACKEND_CYCLE</code>	Unstalled Back-end Cycles on page 108.

`MEMORY_CYCLE` counts the number of cycles that the pipeline is stalled when instructions are waiting for source operands from the memory subsystem, and for pipeline flushes related to memory-access (L1D way mispredictions, DTC flushes). It also counts the number of clocks that the pipeline stalls for the Register Stack Engine to spill or fill registers to/from memory. The `DATA_ACCESS_CYCLE` monitor operates similarly, but excludes RSE activity.

`UNSTALLED_BACKEND_CYCLE` counts the number of cycles that the back-end is processing instructions without delay and the decoupling buffer between the front-end and back-end is empty. In this situation, any effect on the front-end will appear at the back-end of the pipeline. Thus, this monitor reflects the number of cycles where there are no back-end stalls or flushes, and the decoupling buffer is empty, regardless of whether the L1I and ITLB are being hit or missed. The `INST_ACCESS_CYCLE` monitor includes those cycles where there are no back-end stalls or flushes, the decoupling buffer is empty, and the front-end is stalled waiting on an L1I or ITLB miss.

`PIPELINE_ALL_FLUSH_CYCLE` counts the number of cycles lost to branch related restees. Restees can be classified as branch prediction restees (which occur when the front-end correctly predicts a taken branch) or as branch misprediction restees (which occur when the back-end determines that the front-end incorrectly predicted a taken or not-taken branch). Note that taken branches incorrectly predicted by the front-end will not be counted twice. The branch misprediction flush that occurs in the back-end will override the front-end bubble. The monitor also counts ALAT flushes, serialization flushes, MMU-IEU bypass flushes, failed control speculation flushes and other exception flushes. The monitor `PIPELINE_BACKEND_FLUSH_CYCLE` operates similarly, but excludes front-end restees to correctly predicted branches (commonly known as “branch bubbles”).

[Table 7-11](#) defines derived stall cycle accounting monitors in terms of directly measured monitors.

**Table 7-11. Stall Cycle Monitors (Derived)**

Stall Cycle Monitors (Derived)	Description	Performance Monitor Equation
RSE_ACTIVE_CYCLE.d	RSE Active Cycles	MEMORY_CYCLE - DATA_ACCESS_CYCLE
ISSUE_LIMIT_CYCLE.d	Issue Limit Cycles	DEPENDENCY_ALL_CYCLE - DEPENDENCY_SCOREBOARD_CYCLE
TAKEN_BRANCH_CYCLE.d	Taken Branch Cycles	PIPELINE_ALL_FLUSH_CYCLE - PIPELINE_BACKEND_FLUSH_CYCLE
UNSTALLED_PIPELINE_CYCLE.d	Unstalled Pipeline Cycles	UNSTALLED_BACKEND_CYCLE - INST_ACCESS_CYCLE

## 7.5 Branch Events

The five measured Itanium processor branch events listed in [Table 7-12](#) expand into over fifty measurable branch metrics by using the unit masks described on the event pages. BRANCH\_PATH provides insight into the accuracy of taken/not-taken predicate predictions; unit masks allow classification by prediction, outcome and predictor type. BRANCH\_PREDICTOR classifies how branches are predicted by different predictors as they move down the branch prediction pipeline; unit masks provide finer resolution and break down events into correct predictions, incorrect predicate predictions, and incorrect target predictions. BRANCH\_MULTIWAY collects events exclusively for predictions on multiway branch bundles, from which their single-way counterparts can be derived. BRANCH\_TAKEN\_SLOT gives information regarding the position within a bundle that actually-taken branches occupy. BRANCH\_EVENT counts the number of events captured in the branch trace buffer.

[Table 7-13](#) defines derived branch monitors in terms of directly measure monitors.

**Table 7-12. Branch Monitors**

Branch Events	Description
BRANCH_PATH	Accuracy of predicate (taken/not-taken) predictions.
BRANCH_PREDICTOR	Classification of how the branches are predicted in the pipeline.
BRANCH_MULTIWAY	Details on multiway branch bundle predictions (details on single-way branch bundle predictions can be derived from this event).
BRANCH_TAKEN_SLOT	Location of taken branches (if any) in a bundle.
BRANCH_EVENT	Branch Event Captured.

**Table 7-13. Branch Monitors (Derived)**

Branch Events	Description	Performance Monitor Equation
BRANCH_MISPREDICTIONS.d	Branch Bundles Mispredicted	(BRANCH_PREDICTOR.ALL.ALL_PREDICTIONS - BRANCH_PREDICTOR.ALL.CORRECT_PREDICTIONS) or (BRANCH_PREDICTOR.ALL.WRONG_PATH + BRANCH_PREDICTOR.ALL.WRONG_TARGET)
BRANCH_1ST_STAGE_PREDICTIONS.d	Branch Bundles (Correctly or Incorrectly) Predicted in the 1st Pipeline Stage	BRANCH_PREDICTOR.1ST_STAGE.ALL_PREDICTIONS
BRANCH_1ST_STAGE_MISPREDICTIONS.d	Branch Bundles Incorrectly Predicted in the 1st Pipeline Stage	(BRANCH_PREDICTOR.ALL.ALL_PREDICTIONS - BRANCH_PREDICTOR.ALL.CORRECT_PREDICTIONS) or (BRANCH_PREDICTOR.1ST_STAGE.WRONG_PATH + BRANCH_PREDICTOR.1ST_STAGE.WRONG_TARGET)
BRANCH_2ND_STAGE_PREDICTIONS.d	Branch Bundles (Correctly or Incorrectly) Predicted in the 2nd Pipeline Stage	BRANCH_PREDICTOR.2ND_STAGE.ALL_PREDICTIONS
BRANCH_2ND_STAGE_MISPREDICTIONS.d	Branch Bundles Incorrectly Predicted in the 2nd Pipeline Stage	(BRANCH_PREDICTOR.ALL.ALL_PREDICTIONS - BRANCH_PREDICTOR.ALL.CORRECT_PREDICTIONS) or (BRANCH_PREDICTOR.2ND_STAGE.WRONG_PATH + BRANCH_PREDICTOR.2ND_STAGE.WRONG_TARGET)
BRANCH_3RD_STAGE_PREDICTIONS.d	Branch Bundles (Correctly or Incorrectly) Predicted in the 3rd Pipeline Stage	BRANCH_PREDICTOR.3RD_STAGE.ALL_PREDICTIONS
BRANCH_3RD_STAGE_MISPREDICTIONS.d	Branch Bundles Incorrectly Predicted in the 3rd Pipeline Stage	(BRANCH_PREDICTOR.ALL.ALL_PREDICTIONS - BRANCH_PREDICTOR.ALL.CORRECT_PREDICTIONS) or (BRANCH_PREDICTOR.3RD_STAGE.WRONG_PATH + BRANCH_PREDICTOR.3RD_STAGE.WRONG_TARGET)
BRANCH_MULTIWAY_COMPONENT.d	Multiway Branch Bundle Predictions Relative to All Prediction	BRANCH_MULTIWAY.ALL_PATHS.ALL_PREDICTIONS / BRANCH_PREDICTOR.ALL.ALL_PREDICTIONS

All branch events can be qualified by instruction address range and opcode matching as described in [Section 6.1.3, “Event Qualification”](#). Since the instruction address range check is bundle granular, qualification of multiway branches by address range is straightforward. However, for opcode matching purposes, multiway branches (MBB or BBB bundle templates) are qualified up to and including the first taken branch as follows:

```
((address range and opcode match on instruction slot 0)
  and (branch in slot 0 is taken))
or ((address range and opcode match on instruction slot 1)
  and (branch in slot 0 is NOT taken)
  and (branch in slot 1 is taken))
or ((address range and opcode match on instruction slot 0 or 1 or 2)
  and (branch in slot 0 is NOT taken)
  and (branch in slot 1 is NOT taken))
```

## 7.6 Memory Hierarchy

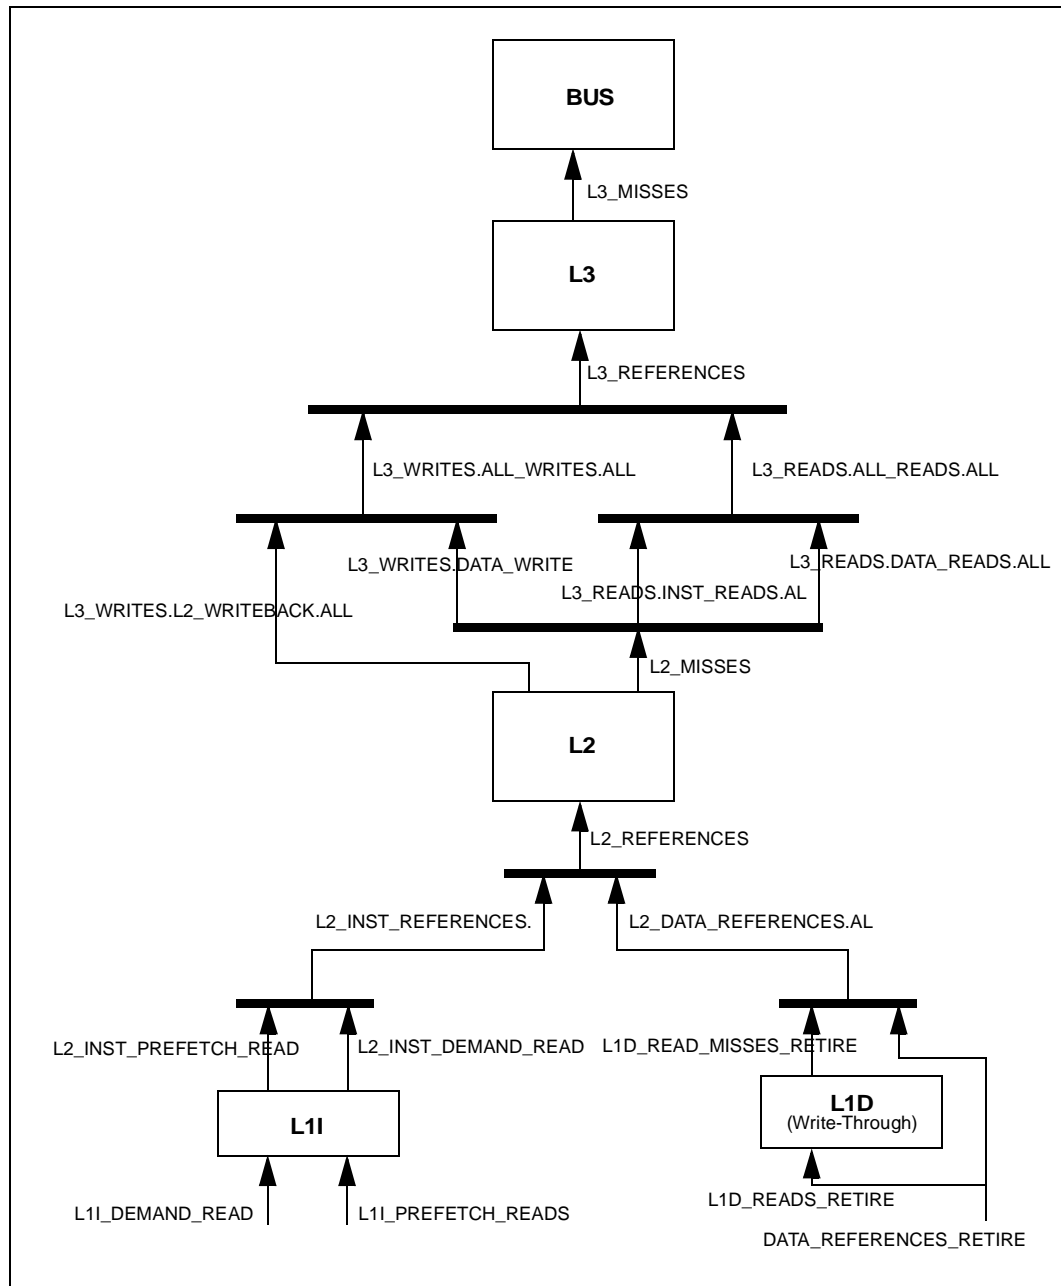
This section summarizes events related to the Itanium processor's memory hierarchy. The memory hierarchy events are grouped as follows:

- L1 Instruction Cache and Prefetch ([Section 7.6.1](#))
- L1 Data Cache ([Section 7.6.2](#))
- L2 Unified Cache ([Section 7.6.3](#))
- L3 Unified Cache ([Section 7.6.4](#))

An overview of the Itanium processor's three-level memory hierarchy and its event monitors is shown in [Figure 7-1](#). The instruction and the data stream work through separate L1 caches. The L1 data cache is a write-through cache. A unified L2 cache serves both the L1 instruction and data caches, and is backed by a large unified L3 cache. Events for individual levels of the cache hierarchy are described in the following three sections. They can be used to compute the most common cache performance ratios summarized in [Table 7-15](#).

For common performance metrics not directly measured by hardware, the equations listed in [Table 7-14](#) can be used.

Figure 7-1. Event Monitors in the Intel® Itanium™ Processor Memory Hierarchy



**Table 7-14. Derived Memory Hierarchy Monitors**

Memory Hierarchy Monitors (Derived)	Description	Performance Monitor Equation
L1I_REFERENCES.d	L1 Instruction Cache References	L1I_PREFETCH_READS + L1I_DEMAND_READS
L2_INST_REFERENCES.d	L2 Instruction References	L2_INST_DEMAND_READS + L2_INST_PREFETCH_READS
L3_DATA_REFERENCES.d	L3 Data References	L3_WRITES.DATA_WRITES.ALL+ L3_READS.DATA_READS.ALL
L3_CORRECTION_RATIO.d	L3 Correction Ratio	L2_MISSES/(L3_REFERENCES-L3_WRITES.L2_WRITEBACK.ALL)

Because Itanium performance monitors for the L2 cache contain secondary misses from L1, and performance monitors for L3 do not, one cannot directly compare L2 and L3 performance monitors. In the performance metrics that follow in [Table 7-15](#) where possible only events from an L2 view or an L3 view are used in computing the metric. However, several useful metrics cannot be calculated accurately this way. For those metrics we compute a correction ratio to scale L3 events to approximate a value that contains secondary misses. [Table 7-14](#) contains the definition of this correction ratio.

**Table 7-15. Cache Performance Ratios**

Performance Metric	Performance Monitor Equation
L1I Miss Ratio	L2_INST_REFERENCES.d/L1I_REFERENCES.d
L1I Demand Miss Ratio	L2_INST_DEMAND_READS / L1I_DEMAND_READS
L1I Prefetch Miss Ratio	L2_INST_PREFETCH_READS / L1I_PREFETCH_READS
L1D Read Miss Ratio	L1D_READ_MISSES_RETIRED / L1D_READS_RETIRED
L2 Miss Ratio	L2_MISSES / L2_REFERENCES
Approximate L2 Data Miss Ratio	(L3_DATA_REFERENCES.d / L2_DATA_REFERENCES.ALL)*L3_CORRECTION_RATIO.d
Approximate L2 Instruction Miss Ratio (includes prefetches)	(L3_READS.INST_READS.ALL / L2_INST_REFERENCES.d)*L3_CORRECTION_RATIO.d
Approximate L2 Data Read Miss Ratio	(L3_READS.DATA_READS.ALL / L2_DATA_REFERENCES.READS) * L3_CORRECTION_RATIO.d
Approximate L2 Data Write Miss Ratio	(L3_WRITES.DATA_WRITES.ALL / L2_DATA_REFERENCES.WRITES)*L3_CORRECTION_RATIO.d
L2 Instruction Ratio	L2_INST_REFERENCES.d/ L2_REFERENCES
L2 Data Ratio	L2_DATA_REFERENCES.ALL / L2_REFERENCES
L3 Miss Ratio	L3_MISSES / (L3_REFERENCES-L3_WRITES.L2_WRITEBACK.ALL)
L3 Data Miss Ratio	(L3_READS.DATA_READS.MISS + L3_WRITES.DATA_WRITES.MISS) / L3_DATA_REFERENCES.d
L3 Instruction Miss Ratio	L3_READS.INST_READS.MISS / L3_READS.INST_READS.ALL
L3 Data Read Ratio	L3_READS.DATA_READS.ALL / L3_DATA_REFERENCES.d
L3 Data Ratio	L3_DATA_REFERENCES.d / L3_REFERENCES
L3 Instruction Ratio	L3_READS.INST_READS.ALL / L3_REFERENCES



## 7.6.1 L1 Instruction Cache and Prefetch

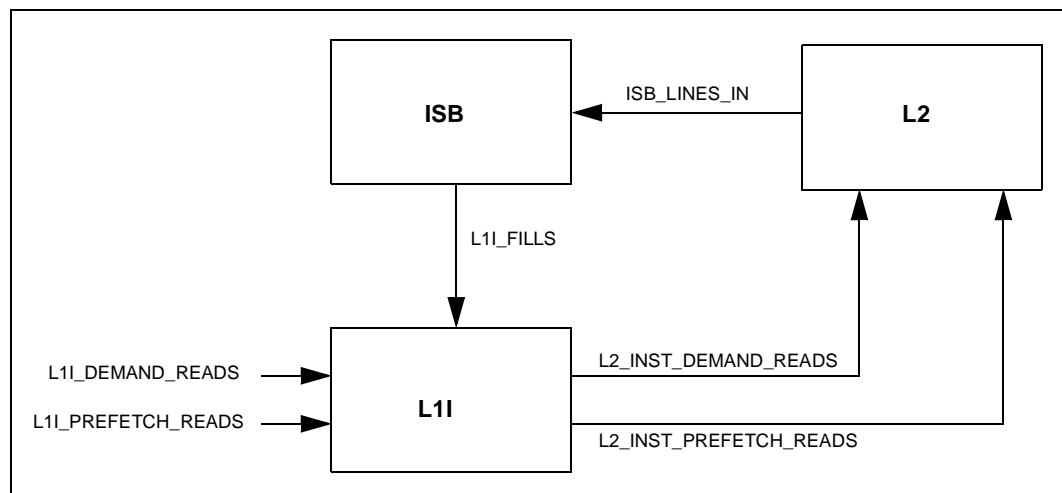
Table 7-16 and Figure 7-2 describes and summarizes the events that the Itanium processor provides to monitor the L1 instruction cache demand fetch and prefetch activity. Table 7-14 lists pertinent derived events. The instruction fetch monitors distinguish between demand fetch (L1I\_DEMAND\_READS, L2\_INST\_DEMAND\_READS) and prefetch activity (L1I\_PREFETCH\_READS, L2\_INST\_PREFETCH\_READS). The amount of data returned from the L2 into the L1 instruction cache and the Instruction Streaming Buffer is monitored by two events (L1I\_FILLS, ISB\_LINES\_IN). The INSTRUCTION\_EAR\_EVENTS monitor (not shown in Figure 7-2) counts how many instruction cache or ITLB misses are captured by the instruction event address register.

The L1 instruction cache and prefetch events can be qualified by the instruction address range check, but not by the opcode matcher. Since instruction cache and prefetch events occur early in the processor pipeline, they include events caused by speculative, wrong-path as well as predicated off instructions. Since the address range check is not based on actually retired, but speculative instruction addresses, event counts may be inaccurate when the range checker is confined to address ranges smaller than the length of the processor pipeline (see Section 6.2.4, “Intel® Itanium™ Instruction Address Range Check Register (PMC[13])” for details).

**Table 7-16. L1 Instruction Cache and Instruction Prefetch Monitors**

L1I and I-Prefetch Monitors	Description
L1I_DEMAND_READS	L1I and ISB Instruction Demand Lookups on page 96.
L1I_FILLS	L1 Instruction Cache Fills on page 96.
L2_INST_DEMAND_READS	L2 Instruction Demand Fetch Requests on page 98.
INSTRUCTION_EAR_EVENTS	Instruction EAR Events on page 94.
L1I_PREFETCH_READS	L1I and ISB Instruction Prefetch Lookups on page 97.
L2_INST_PREFETCH_READS	L2 Instruction Prefetch Requests on page 99.
ISB_LINES_IN	Instruction Streaming Buffer Lines In on page 95.

**Figure 7-2. L1 Instruction Cache and Prefetch Monitors**



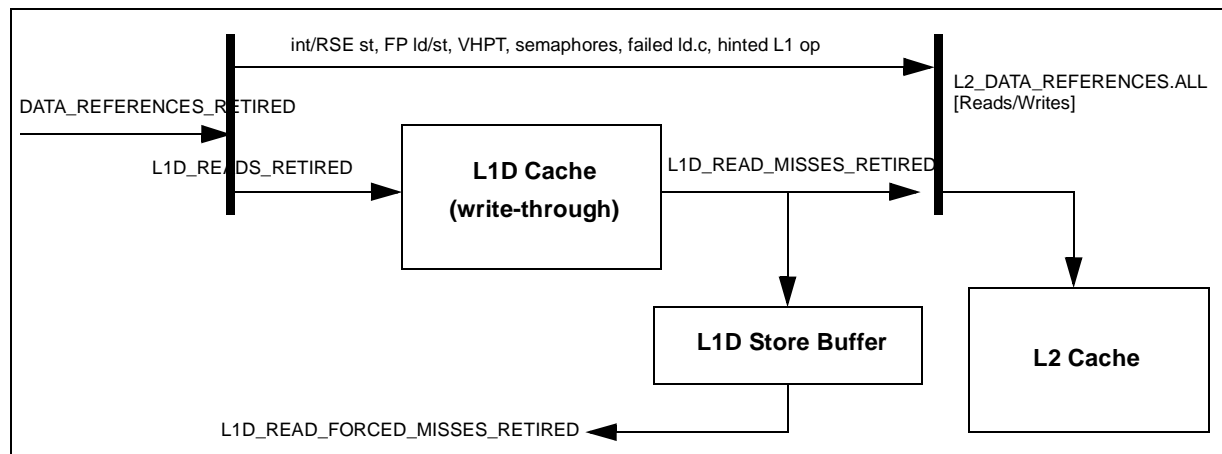
## 7.6.2 L1 Data Cache

Table 7-17 lists the Itanium processor's seven L1 data cache monitors. As shown in Figure 7-3, the write-through L1 data cache services cacheable loads, Integer and RSE stores, FP memory operations, VHPT references, semaphores, check loads and hinted L2 memory references. DATA\_REFERENCES\_RETIRED is the number of issued data memory references. L1 data cache reads (L1D\_READS\_RETIRED) and L1 data cache misses (L1D\_READ\_MISSES\_RETIRED) monitor the read hit/miss rate for the L1 data cache. The number of L2 data references (L2\_DATA\_REFERENCES.ALL) is the number of data requests prior to cache line merging. Unit mask selections allow breaking down into reads and writes. The DATA\_EAR\_EVENTS monitor (not shown in Figure 7-3) counts how many data cache or DTLB misses are captured by the data event address register. RSE operations are included in all data cache monitors, but are not broken down explicitly.

**Table 7-17. L1 Data Cache Monitors**

L1D Monitors	Description
DATA_REFERENCES_RETIRED	Retired Data Memory References on page 89.
L1D_READS_RETIRED	L1 Data Cache Reads on page 96.
L1D_READ_MISSES_RETIRED	L1 Data Cache Read Misses on page 96.
PIPELINE_FLUSH.L1D_WAY_MISPREDICT	Pipeline Flush on page 106.
L1D_READ_FORCED_MISSES_RETIRED	L1 Data Cache Forced Load Misses on page 95.
L2_DATA_REFERENCES.ALL	L2 Data Read and Write References on page 97.
DATA_EAR_EVENTS	L1 Data Cache EAR Events on page 89.

**Figure 7-3. L1 Data Cache Monitors**



## 7.6.3 L2 Unified Cache

Table 7-18 summarizes the directly-measured events that monitor the Itanium processor L2 cache. Table 7-14 lists pertinent derived events. Refer to Figure 7-1 for a graphical view of the L2 cache monitors.

L2\_REFERENCES, L2\_INST\_DEMAND\_READS, L2\_INST\_PREFETCH\_READS, L2\_DATA\_REFERENCES.ALL, and L2\_MISSES are all counted in terms of number of requests seen by the L2. L2\_FLUSHES and L2\_FLUSH\_DETAILS count and break-down the number of L2 flushes due to address conflicts, store buffer conflicts, bus rejects, and other reasons. L1D\_READ\_FORCED\_MISSES\_RETIRED counts the number of loads that were bypassed from an earlier store.

**Table 7-18. L2 Cache Monitors**

L1 Monitors	Description
L2_REFERENCES	L2 References on page 99.
L2_INST_PREFETCH_READS	L2 Instruction Prefetch Requests on page 99.
L2_INST_DEMAND_READS	L2 Instruction Demand Fetch Requests on page 98.
L2_DATA_REFERENCES.ALL	L2 Data Read and Write References on page 97.
L2_DATA_REFERENCES.READS	L2 Data Read References on page 97.
L2_DATA_REFERENCES.WRITEES	L2 Data Write References on page 97.
L2_MISSES	L2 Misses on page 99.
L2_FLUSHES	L2 Flushes on page 98.
L2_FLUSH_DETAILS	L2 Flush Details on page 98.

## 7.6.4 L3 Unified Cache

Table 7-19 summarizes the directly-measured L3 cache events. Table 7-14 lists pertinent derived events. Refer to Figure 7-1 for a graphical view of the L3 cache monitors.

**Table 7-19. L3 Cache Monitors**

L2 Monitors	Description
L3_REFERENCES	L3 References on page 102.
L3_MISSES	L3 Misses on page 100.
L3_LINES_REPLACED	L3 Cache Lines Replaced on page 99.
L3_READS.ALL_READS.ALL	Instruction and Data L3 Reads on page 100.
L3_READS.ALL_READS.HIT	Instruction and Data L3 Read Hits on page 100.
L3_READS.ALL_READS.MISS	Instruction and Data L3 Read Misses on page 100.
L3_READS.DATA_READS.ALL	Data L3 Reads on page 100.
L3_READS.DATA_READS.HIT	Data L3 Read Hits on page 101.
L3_READS.DATA_READS.MISS	Data L3 Read Misses on page 101.
L3_READS.INST_READS.ALL	Instruction L3 Reads on page 101.
L3_READS.INST_READS.HIT	Instruction L3 Read Hits on page 101.
L3_READS.INST_READS.MISS	Instruction L3 Read Misses on page 101.
L3_WRITES.ALL_WRITES.ALL	L3 Writes on page 102.
L3_WRITES.ALL_WRITES.HIT	L3 Write Hits on page 102.

**Table 7-19. L3 Cache Monitors (Continued)**

L2 Monitors	Description
L3_WRITES.ALL_WRITES.MISS	L3 Write Misses on page 102.
L3_WRITES.L2_WRITEBACK.ALL	L3 Writebacks on page 103.
L3_WRITES.L2_WRITEBACK.HIT	L3 Writeback Hits on page 103.
L3_WRITES.L2_WRITEBACK.MISS	L3 Writeback Misses on page 103.
L3_WRITES.DATA_WRITES.ALL	L3 Data Writes on page 103.
L3_WRITES.DATA_WRITES.HIT	L3 Data Write Hits on page 103.
L3_WRITES.DATA_WRITES.MISS	L3 Data Write Misses on page 104.

## 7.6.5 Frontside Bus

Table 7-20 lists the frontside bus or system bus transaction monitors.

**Table 7-20. Bus Events**

L2 Monitors	Description
BUS_ALL	Bus Transactions on page 80.
BUS_PARTIAL	Bus Partial Transactions on page 84.
BUS_BURST	Bus Burst Transactions on page 82.
BUS_MEMORY	Bus Memory Transactions on page 84.
BUS_RD_ALL	Bus Read Transactions on page 84.
BUS_RD_DATA	Bus Read Data Transactions on page 84.
BUS_RD_PRTL	Bus Read Partial Transactions on page 86.
BUS_RD_HIT	Bus Read Hit Clean Non-local Cache Transactions on page 85.
BUS_RD_HITM	Bus Read Hit Modified Non-local Cache Transactions on page 85.
BUS_RD_INVAL	Bus Read Invalidated Line on page 85.
BUS_RD_INVAL_HITM	Bus BIL Transaction Results in HITM on page 86.
BUS_RD_INVAL_BST	Bus BRIL Burst Transactions on page 85.
BUS_RD_INVAL_BST_HITM	Bus BRIL Burst Transaction Results in HITM on page 86.
BUS_HITM	Bus Hit Modified Line Transactions on page 82.
BUS_WR_WB	Bus Write Back Transactions on page 88.
BUS_SNOOPS_HITM	Bus Snoops Hit Modified Cache Line on page 87.
BUS_SNOOPS	Bus Snoops Total on page 87.
BUS_SNOOP_STALL_CYCLES	Bus Snoop Stall Cycles on page 87.
BUS_SNOOPQ_REQ	Bus Snoop Queue Requests, Category: Frontside Bus on page 88.
BUS_BRQ_READ_REQ_INSERTED	BRQ Requests Inserted on page 81.
BUS_IO	IA-32 Compatible I/O Bus Transactions on page 82.
BUS_RD_IO	IA-32 Compatible I/O Read Transactions on page 86.
BUS_LOCK	IA-32 Compatible Bus Lock Transactions on page 83.
BUS_LOCK_CYCLES	IA-32 Compatible Bus Lock Cycles on page 83.

Table 7-21 lists the derived frontside bus transaction monitors.

**Table 7-21. Frontside Bus Monitors (Derived)**

Bus Monitors (Derived)	Description	Performance Monitor Equation
BUS_RD_INSTRUCTIONS.d	Bus Read Instructions	BUS_RD_ALL - BUS_RD_DATA
BUS_RD_INVAL_MEMORY.d	Bus BIL Transaction Satisfied from Memory	BUS_RD_INVAL - BUS_RD_INVAL_HITM
BUS_RD_INVAL_BST_MEMORY.d	Bus BRIL Burst Transaction Satisfied from Memory	BUS_RD_INVAL_BST - BUS_RD_INVAL_BST_HITM
BUS_ADDR_BPRI.d	Bus Used by I/O Agent	BUS_MEMORY.IOAGENT
BUS_IOQ_LIVE_REQ.d	In-order Bus Queue Requests	BUS_IOQ_LIVE_REQ_HI * 4 + BUS_IOQ_LIVE_REQ_LO
BUS_BRQ_LIVE_REQ.d	BRQ Live Requests	BUS_BRQ_LIVE_REQ_HI * 4 + BUS_BRQ_LIVE_REQ_LO

Most of the bus events in Section 7.6.5 can be qualified by the bus transaction initiator using the three way unit mask as described in Table 7-22.

**Table 7-22. Unit Masks for Qualifying Bus Transaction Events by Initiator**

Selection	PMC.umask[19:16]	Description
ANY	x001	Counts all bus transactions (initiated by any processor or non-processor bus masters)
SELF	x010	Counts bus transactions initiated by the local processor only
IO	x100	Counts bus transactions from IO agents, i.e. non-processor bus masters

Table 7-23 defines the conventions that will be used when describing the Itanium processor frontside bus transaction monitors in Section 7.6.5.

**Table 7-23. Conventions for Frontside Bus Transactions**

Name	Description
BRL	Memory Read (64 byte bursts). Includes code fetches and data loads from WB memory.
BRIL	Memory Read & Invalidate (64 byte bursts). Also known as read for ownership (RFO).
BIL	Memory Read & Invalidate (0 byte sized transaction). Caused by flush cache (fc) instruction only.
BWL	Memory Write (64 byte bursts). Explicit writebacks/coalesced writes.
BRP	Partial Memory Reads (<64 byte transactions). Typically, uncacheable reads.
BWP	Partial Memory Write (<64 byte transactions). Typically, uncacheable writes.
IORD	Partial IO Read (<64 byte transactions). Uncacheable read to IO port space.
IOWR	Partial IO Write (<64 byte transactions). Uncacheable write to IO port space.

Other transactions besides those listed in Table 7-23 include Deferred Reply, Special Transactions, Interrupt, Interrupt Acknowledge, and Purge TC. For the bus performance monitors in Section 7.6.5, note that the monitors will count if any transaction gets a retry response from the priority agent.

To support the analysis of snoop traffic in a multiprocessor system, the Itanium processor provides local processor and remote response monitors. The local processor snoop events (BUS\_SNOOPS\_HITM, BUS\_SNOOPS, BUS\_SNOOPQ\_REQ) monitor inbound snoop traffic. The remote response events (BUS\_RD\_HIT, BUS\_RD\_HITM, BUS\_RD\_INVALID\_HITM, BUS\_RD\_INVALID\_BST\_HITM) monitor the snoop responses of other processors to bus transactions that the monitoring processor originated. [Table 7-24](#) summarizes the remote snoop events by bus transaction.

**Table 7-24. Bus Events by Snoop Response**

Remote Processor Response	BRL	BIL	BRIL
HIT	BUS_RD_HIT	NA	NA
HITM	BUS_RD_HITM	BUS_RD_INVALID_HITM	BUS_RD_INVALID_BST_HITM
ALL	BUS_RD_ALL	BUS_RD_INVALID	BUS_RD_INVALID

With the Itanium processor frontside bus monitors, the performance metrics described in [Table 7-25](#) can be computed.

**Table 7-25. Bus Performance Metrics**

Performance Metric	Performance Monitor Equation
Cacheable Data Fetch Bus Transaction Ratio	$\text{BUS\_RD\_DATA}/\text{BUS\_ALL}$ or $\text{BUS\_RD\_DATA}/\text{BUS\_MEMORY}$
Partial Access Ratio	$\text{BUS\_PARTIAL}/\text{BUS\_MEMORY}$
Read Partial Access Ratio	$\text{BUS\_RD\_PRTL}/\text{BUS\_MEMORY}$
Read Hit To Shared Line Ratio	$\text{BUS\_RD\_HIT}/\text{BUS\_RD\_ALL}$ or $\text{BUS\_MEMORY}$
Read Hit to Modified Line Ratio	$\text{BUS\_RD\_HITM}/\text{BUS\_RD\_ALL}$ or $\text{BUS\_RD\_HITM}/\text{BUS\_MEMORY}$
BIL Ratio	$\text{BUS\_RD\_HIT}/\text{BUS\_MEMORY}$
BIL Hit to Modified Line Ratio	$\text{BUS\_RD\_INVALID\_HITM}/\text{BUS\_MEMORY}$ or $\text{BUS\_RD\_INVALID\_HITM}/\text{BUS\_RD\_INVALID}$
BRIL Hit to Modified Line Ratio	$\text{BUS\_RD\_INVALID\_BST\_HITM}/\text{BUS\_MEMORY}$ or $\text{BUS\_RD\_INVALID\_BST\_HITM}/\text{BUS\_RD\_INVALID}$
Bus Modified Line Hit Ratio	$\text{BUS\_RD\_HITM}/\text{BUS\_MEMORY}$ or $\text{BUS\_RD\_HITM}/\text{BUS\_BURST}$
Writeback Ratio	$\text{BUS\_WR\_WB}/\text{BUS\_MEMORY}$ or $\text{BUS\_WR\_WB}/\text{BUS\_BURST}$
Cacheable Read Ratio	$(\text{BUS\_RD\_ALL} + \text{BUS\_RD\_INVALID\_BST})/\text{BUS\_MEMORY}$
I/O Cycle Ratio	$\text{BUS\_IO}/\text{BUS\_ALL}$
I/O Read	$\text{BUS\_RD\_IO}/\text{BUS\_ALL}$

## 7.7 System Events

Table 7-26 lists the directly measurable system and TLB events. Table 7-27 lists pertinent derived events. The debug register match events count how often the address in any instruction or data break-point register (IBR or DBR) matches the current retired instruction pointer (CODE\_DEBUG\_REGISTER\_MATCHES.d) or the current data memory address (DATA\_DEBUG\_REGISTER\_MATCHES.d). PIPELINE\_FLUSH counts the number of times the Itanium processor pipeline is flushed due to a data translation cache miss, L1 data cache way mispredict, an exception flush or an instruction serialization event. CPU\_CPL\_CHANGES counts the number of privilege level transitions due to interruptions, system calls (epc) and returns (demoting branch), and `r f i` instructions. CPU\_CYCLES counts the number of cycles the CPU is not powered down or in light HALT state.

**Table 7-26. System and TLB Monitors**

System and Processor TLB Monitors	Description
PIPELINE_FLUSH	Pipeline Flush on page 106.
CPU_CPL_CHANGES	Privilege level changes on page 88.
CPU_CYCLES	CPU Cycles on page 88.
ITLB_MISSES_FETCH	ITLB Demand Misses on page 95.
ITLB_INSERTS_HPW	Hardware Page Walker Inserts into the ITLB on page 95.
DTC_MISSES	DTC Misses on page 90.
DTLB_MISSES	DTLB Misses on page 90.
DTLB_INSERTS_HPW	Hardware Page Walker Inserts into the DTLB on page 90.

Table 7-27 defines derived system and TLB events that are computed from events directly measured by hardware.

**Table 7-27. System and TLB Monitors (Derived)**

Derived Memory Hierarchy Monitors	Description	Performance Monitor Equation
CODE_DEBUG_REGISTER_MATCHES.d	Code Debug Register Matches	IA64_TAGGED_INST_RETIRED
DATA_DEBUG_REGISTER_MATCHES.d	Data Debug Register Matches	LOADS_RETIRED + STORES_RETIRED
ITLB_REFERENCES.d	ITLB References	L1I_DEMAND_READS
ITLB_EAR_EVENT.d	ITLB EAR Event	INSTRUCTION_EAR_EVENTS
DTLB_REFERENCES.d	DTLB References	DATA_REFERENCES_RETIRED
DTLB_EAR_EVENT.d	DTLB EAR Event	DATA_EAR_EVENTS

The Itanium processor instruction and data TLBs and the virtual hash page table walker are monitored by the events described in Table 7-26 and Table 7-27. Figure 7-4 gives a graphical summary. Table 7-28 lists the TLB performance metrics that can be computed using these events.

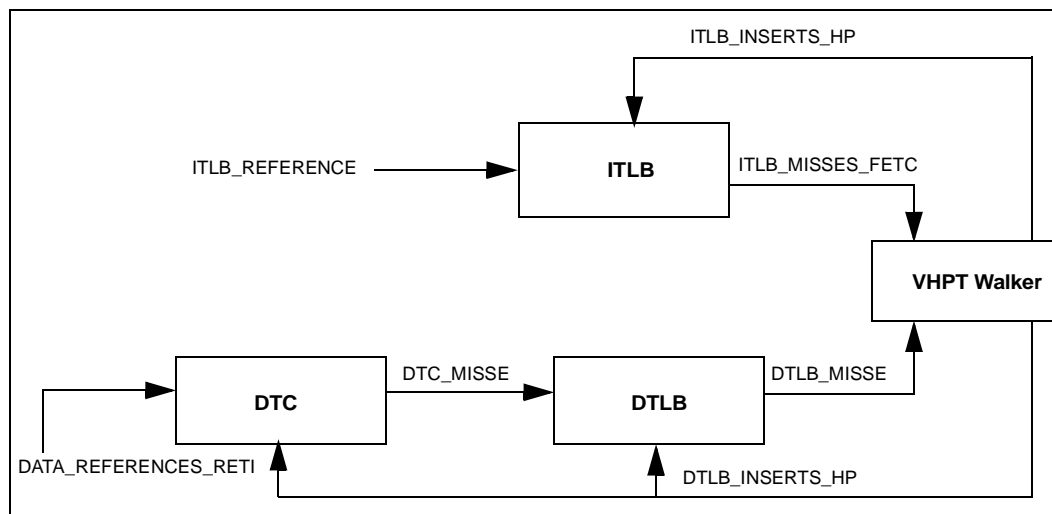
ITLB\_REFERENCES.d and DTLB\_REFERENCES.d are derived from the respective instruction/data cache access events. Note that ITLB\_REFERENCES.d does not include prefetch requests made to the L1I cache (L1I\_PREFETCH\_READS). This is because prefetches are cancelled when they miss in the ITLB and thus do not trigger VHPT walks or software TLB miss handling. ITLB\_MISSES\_FETCH and DTLB\_MISSES count TLB misses.

ITLB\_INSERTS\_HPW and DTLB\_INSERTS\_HPW count the number of instruction/data TLB inserts performed by the virtual hash page table walker. The Itanium processor data TLB is a two level TLB; DTC\_MISSES counts the number of first level data TLB misses.

**Table 7-28. TLB Performance Metrics**

Performance Metric	Performance Monitor Equation
ITLB Miss Ratio	$ITLB\_MISSES\_FETCH / ITLB\_REFERENCES.d$
DTLB Miss Ratio	$DTLB\_MISSES / DTLB\_REFERENCES.d$
DTC Miss Ratio	$DTC\_MISSES / DTLB\_REFERENCES.d$

**Figure 7-4. Instruction and Data TLB Monitors**



## 7.8 Performance Monitor Event List

This section enumerates Itanium processor performance monitoring events.

### ALAT\_REPLACEMENT.ALL

- **Title:** ALAT Entries Replaced by Any Instruction, **Category:** Execution
- **Definition:** ALAT\_REPLACEMENT.ALL counts the number of times an advanced load (`ld.a` or `ld.as` or `ldfp.a` or `ldfp.as`) or a no-clear check load (`ld.c.nc` and variants of `ldf.c.nc`) displaced a valid entry in the ALAT
- **Event Code:** 0x38, **Umask:** xx11, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes



### ALAT\_REPLACEMENT.FP

- **Title:** ALAT Entries Replaced by FP Instructions, **Category:** Execution
- **Definition:** ALAT\_REPLACEMENT.FP counts the number of times a FP advanced load (`ldfp.a` or `ldfp.as`) or a no-clear FP check load (variants of `ldf.c.nc`) displaced a valid entry in the ALAT
- **Event Code:** 0x38, **Umask:** xx10, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

### ALAT\_REPLACEMENT.INTEGER

- **Title:** ALAT Entries Replaced by Integer Instructions, **Category:** Execution
- **Definition:** ALAT\_REPLACEMENT.INTEGER counts the number of times an integer advanced load (`ld.a` or `ld.as`) or a no-clear integer check load (`ld.c.nc`) displaced a valid entry in the ALAT
- **Event Code:** 0x38, **Umask:** xx01, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

### ALAT\_INST\_CHKA\_LDC.ALL

- **Title:** Advanced Load Checks and Check Loads, **Category:** Execution
- **Definition:** ALAT\_INST\_CHKA\_LDC.ALL counts the number of all advanced load checks (`chk.a`) and check loads in both clear and no-clear forms (`ld.c.clr` or `ld.c.nc`, including FP variants) as seen by the ALAT
- **Event Code:** 0x36, **Umask:** xx11, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no for `chk.a`, and yes for `ld.c`

### ALAT\_INST\_CHKA\_LDC.FP

- **Title:** FP Advanced Load Checks and Check Loads, **Category:** Execution
- **Definition:** ALAT\_INST\_CHKA\_LDC.FP counts all FP advanced load checks (`chk.a`) and all FP check loads in both clear and no-clear forms (`ld.c.clr` or `ld.c.nc`, FP variants only) as seen by the ALAT
- **Event Code:** 0x36, **Umask:** xx10, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no for `chk.a`, and yes for `ld.c`

**ALAT\_INST\_CHKA\_LDC.INTEGER**

- **Title:** Integer Advanced Load Checks and Check Loads, **Category:** Execution
- **Definition:** ALAT\_INST\_CHKA\_LDC.INTEGER counts all integer advanced load checks (chk . a) and all integer check loads in both clear and no-clear forms (ld . c . clr or ld . c . nc, excluding FP variants) as seen by the ALAT
- **Event Code:** 0x36, **Umask:** xx01, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no for chk . a, and yes for ld . c

**ALAT\_INST\_FAILED\_CHKA\_LDC.ALL**

- **Title:** Failed Advanced Load Checks and Check Loads, **Category:** Execution
- **Definition:** ALAT\_INST\_FAILED\_CHKA\_LDC.ALL counts failed advanced load checks (chk . a) and failed check loads in both clear and no-clear forms (ld . c . clr or ld . c . nc, including FP variants) as seen by the ALAT
- **Event Code:** 0x37, **Umask:** xx11, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no for chk . a, and yes for ld . c

**ALAT\_INST\_FAILED\_CHKA\_LDC.FP**

- **Title:** Failed FP Advanced Load Checks and Check Loads, **Category:** Execution
- **Definition:** ALAT\_INST\_FAILED\_CHKA\_LDC.FP counts failed FP advanced load checks (chk . a) and failed FP check loads in both clear and no-clear forms (ld . c . clr or ld . c . nc, FP variants only) as seen by the ALAT
- **Event Code:** 0x37, **Umask:** xx10, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no for chk . a, and yes for ld . c

**ALAT\_INST\_FAILED\_CHKA\_LDC.INTEGER**

- **Title:** Failed Integer Advanced Load Checks and Check Loads, **Category:** Execution
- **Definition:** ALAT\_INST\_FAILED\_CHKA\_LDC.INTEGER counts the number of failed integer advanced load checks (chk . a) and failed integer check loads in both clear and no-clear forms (ld . c . clr or ld . c . nc, excluding FP variants) as seen by the ALAT
- **Event Code:** 0x37, **Umask:** xx01, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no for chk . a, and yes for ld . c

## ALL\_STOPS\_DISPERSED

- **Title:** Implicit and Explicit Stops Dispersed, **Category:** Instruction Issue
- **Definition:** ALL\_STOPS\_DISPERSED counts the sum of explicit programmer-specified stops (EXPL\_STOPS\_DISPERSED) and dispersal breaks due to resource limitations and branch instructions (independent of their predicate prediction). The sum includes stops encountered during hardware speculative wrong-path execution (i.e., in the shadow of a flush)
- **Event Code:** 0x2F, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

## BRANCH\_EVENT

- **Title:** Branch Event Captured, **Category:** Branch
- **Definition:** BRANCH\_EVENT counts the number of branch events, including multiway branches captured by the Branch Trace Buffer.
- **Event Code:** 0x11, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

## BRANCH\_MULTIWAY.ALL\_PATHS.ALL\_PREDICTIONS

- **Title:** All Branch Predictions on Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.ALL\_PATHS.ALL\_PREDICTIONS counts all branch predictions made on multiway branch bundles
- **Event Code:** 0x0E, **Umask:** 0000, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

## BRANCH\_MULTIWAY.ALL\_PATHS.CORRECT\_PREDICTIONS

- **Title:** Correct Branch Predictions on Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.ALL\_PATHS.CORRECT\_PREDICTIONS counts all branch predictions on multiway branch bundles that do not necessitate a back-end branch misprediction flush
- **Event Code:** 0x0E, **Umask:** 0001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_MULTIWAY.ALL\_PATHS.WRONG\_PATH**

- **Title:** Incorrect Predicate Predictions on Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.ALL\_PATHS.WRONG\_PATH counts the number of multiway branch bundles whose combined predicate is incorrectly predicted. This includes bundles where all branch instructions are predicted not-taken and any one instruction is actually taken, and those bundles where a branch instruction was predicted taken and either a prior branch instruction in the bundle was actually taken or the predicted instruction was not taken. In any event, the processor restears the front-end to the correct target, i.e., a given multiway bundle can only be mispredicted once
- **Event Code:** 0x0E, **Umask:** 0010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_MULTIWAY.ALL\_PATHS.WRONG\_TARGET**

- **Title:** Incorrect Target Predictions on Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.ALL\_PATHS.WRONG\_TARGET counts the number of multiway branch bundles where a branch instruction is correctly predicted taken, but its target is incorrect
- **Event Code:** 0x0E, **Umask:** 0011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_MULTIWAY.NOT\_TAKEN.ALL\_PREDICTIONS**

- **Title:** All Branch Predictions on Not-Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.NOT\_TAKEN.ALL\_PREDICTIONS is analogous to BRANCH\_MULTIWAY.ALL\_PATHS.ALL\_PREDICTIONS, except it applies only to multiway branch bundles where all branch instructions are actually not taken
- **Event Code:** 0x0E, **Umask:** 1000, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_MULTIWAY.NOT\_TAKEN.CORRECT\_PREDICTIONS**

- **Title:** Correct Branch Predictions on Not-Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.NOT\_TAKEN.CORRECT\_PREDICTIONS is analogous to BRANCH\_MULTIWAY.ALL\_PATHS.CORRECT\_PREDICTIONS, except it applies only to multiway branch bundles where all branch instructions are actually not taken
- **Event Code:** 0x0E, **Umask:** 1001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_MULTIWAY.NOT\_TAKEN.WRONG\_PATH**

- **Title:** Incorrect Predicate Predictions on Not-Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.NOT\_TAKEN.WRONG\_PATH is analogous to BRANCH\_MULTIWAY.ALL\_PATHS.WRONG\_PATH, except it applies only to multiway branch bundles where all branch instructions are actually not taken
- **Event Code:** 0x0E, **Umask:** 1010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_MULTIWAY.NOT\_TAKEN.WRONG\_TARGET**

- **Title:** Incorrect Target Predictions on Not-Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.NOT\_TAKEN.WRONG\_TARGET should always count zero, as not-taken branches do not specify a branch target
- **Event Code:** 0x0E, **Umask:** 1011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_MULTIWAY.TAKEN.ALL\_PREDICTIONS**

- **Title:** All Branch Predictions on Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.TAKEN.ALL\_PREDICTIONS is analogous to BRANCH\_MULTIWAY.ALL\_PATHS.ALL\_PREDICTIONS, except it applies only to multiway branch bundles where at least one branch instruction is taken
- **Event Code:** 0x0E, **Umask:** 1100, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_MULTIWAY.TAKEN.CORRECT\_PREDICTIONS**

- **Title:** Correct Branch Predictions on Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.TAKEN.CORRECT\_PREDICTIONS is analogous to BRANCH\_MULTIWAY.ALL\_PATHS.CORRECT\_PREDICTIONS, except it applies only to multiway branch bundles where at least one branch instruction is taken
- **Event Code:** 0x0E, **Umask:** 1101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_MULTIWAY.TAKEN.WRONG\_PATH**

- **Title:** Incorrect Predicate Predictions on Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.TAKEN.WRONG\_PATH is analogous to BRANCH\_MULTIWAY.ALL\_PATHS.WRONG\_PATH, except it applies only to multiway branch bundles where at least one branch instruction is taken
- **Event Code:** 0x0E, **Umask:** 1110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_MULTIWAY.TAKEN.WRONG\_TARGET**

- **Title:** Incorrect Target Predictions on Taken Multiway Bundles, **Category:** Branch
- **Definition:** BRANCH\_MULTIWAY.TAKEN.WRONG\_TARGET should equal BRANCH\_MULTIWAY.ALL\_PATHS.WRONG\_TARGET, since only multiway branch bundles where at least one branch instruction is taken actually specify a target
- **Event Code:** 0x0E, **Umask:** 1111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.ALL.NT\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Not-Taken Predicate Predictions, **Category:** Branch
- **Definition:** BRANCH\_PATH.ALL.NT\_OUTCOMES\_CORRECTLY\_PREDICTED counts the number of correct not-taken predicate predictions on not-taken branches, independent of predictor
- **Event Code:** 0x0F, **Umask:** 0010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.ALL.TK\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Taken Predicate Predictions, **Category:** Branch
- **Definition:** BRANCH\_PATH.ALL.TK\_OUTCOMES\_CORRECTLY\_PREDICTED counts the number of correct taken predicate predictions on taken branches, independent of predictor. Only the predicate must be correct; resteers to incorrect targets are also counted by this monitor as long as the branch is actually taken
- **Event Code:** 0x0F, **Umask:** 0011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PATH.ALL.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Taken Predicate Predictions, **Category:** Branch
- **Definition:** BRANCH\_PATH.ALL.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED counts the number of incorrect taken predicate predictions on not-taken branches, independent of predictor
- **Event Code:** 0x0F, **Umask:** 0000, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PATH.ALL.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Not-Taken Predicate Predictions, **Category:** Branch
- **Definition:** BRANCH\_PATH.ALL.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED counts the number of incorrect not-taken predicate predictions on taken branches, independent of predictor
- **Event Code:** 0x0F, **Umask:** 0001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PATH.1ST\_STAGE.NT\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Not-Taken Predicate Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.1ST\_STAGE.NT\_OUTCOMES\_CORRECTLY\_PREDICTED should always count zero, as the TAR is the only predictor in the first stage of the core pipeline and it only makes taken predictions
- **Event Code:** 0x0F, **Umask:** 0110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PATH.1ST\_STAGE.TK\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Taken Predicate Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.1ST\_STAGE.TK\_OUTCOMES\_CORRECTLY\_PREDICTED counts the number of correct taken predicate predictions on taken branches made by the TAR in the first stage of the core pipeline. Only the predicate must be correct; restesters to incorrect targets are also counted by this monitor as long as the branch is actually taken. There are 0 bubbles between the branch and its predicted target
- **Event Code:** 0x0F, **Umask:** 0111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.1ST\_STAGE.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Taken Predicate Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.1ST\_STAGE.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED counts the number of incorrect taken predicate predictions on not-taken branches, made by the TAR in the first stage of the core pipeline
- **Event Code:** 0x0F, **Umask:** 0100, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.1ST\_STAGE.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Taken Predicate Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.1ST\_STAGE.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED should always count zero, as the TAR is the only predictor in the first stage of the core pipeline and it only makes taken predictions
- **Event Code:** 0x0F, **Umask:** 0101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.2ND\_STAGE.NT\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Taken Predicate Predictions made in the second pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.2ND\_STAGE.NT\_OUTCOMES\_CORRECTLY\_PREDICTED counts the number of correct not-taken predicate predictions on not-taken branches made by the BPT/MBPT in the second stage of the core pipeline
- **Event Code:** 0x0F, **Umask:** 1010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.2ND\_STAGE.TK\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Taken Predicate Predictions made in the second pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.2ND\_STAGE.TK\_OUTCOMES\_CORRECTLY\_PREDICTED counts the number of correct taken predicate predictions on taken branches by the BPT/MBPT or the TAC in the second stage of the core pipeline. Only the predicate must be correct; restesters to incorrect targets are also counted by this monitor as long as the branch is actually taken. There is 1 bubble between the branch and its predicted target
- **Event Code:** 0x0F, **Umask:** 1011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no



### **BRANCH\_PATH.2ND\_STAGE.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Taken Predicate Predictions made in the second pipeline stage  
**Category:** Branch
- **Definition:** BRANCH\_PATH.2ND\_STAGE.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED counts the number of incorrect taken predicate predictions on not-taken branches made by the BPT/MBPT or the TAC in the second stage of the core pipeline
- **Event Code:** 0x0F, **Umask:** 1000, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PATH.2ND\_STAGE.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Not-Taken Predicate Predictions made in the second pipeline stage,  
**Category:** Branch
- **Definition:** BRANCH\_PATH.2ND\_STAGE.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED counts the number of incorrect not-taken predicate predictions on taken branches made by the BPT/MBPT in the second stage of the core pipeline
- **Event Code:** 0x0F, **Umask:** 1001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PATH.3RD\_STAGE.NT\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Not-Taken Predicate Predictions made in the third pipeline stage,  
**Category:** Branch
- **Definition:** BRANCH\_PATH.3RD\_STAGE.NT\_OUTCOMES\_CORRECTLY\_PREDICTED counts the number of correct not-taken predicate predictions on not-taken branches made by the BAC in the third stage of the core pipeline, including overrides of TAR taken predictions (made in the first stage) on the last instances of loop-closing branches
- **Event Code:** 0x0F, **Umask:** 1110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.3RD\_STAGE.TK\_OUTCOMES\_CORRECTLY\_PREDICTED**

- **Title:** Correct Taken Predicate Predictions made in the third pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.3RD\_STAGE.TK\_OUTCOMES\_CORRECTLY\_PREDICTED counts the number of correct taken predicate predictions on taken branches made by the BAC in the third stage of the core pipeline. Only the predicate must be correct; restesters to incorrect targets are also counted by this monitor as long as the branch is actually taken. There are 2 bubbles between the branch and its predicted target (or 3, if the target must be computed for a branch syllable in slot 0 or 1)
- **Event Code:** 0x0F, **Umask:** 1111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.3RD\_STAGE.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Taken Predicate Predictions made in the third pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.3RD\_STAGE.NT\_OUTCOMES\_INCORRECTLY\_PREDICTED counts the number of incorrect taken predicate predictions on not-taken branches made by the BAC in the third stage of the core pipeline
- **Event Code:** 0x0F, **Umask:** 1100, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PATH.3RD\_STAGE.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED**

- **Title:** Incorrect Not-Taken Predicate Predictions made in the third pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PATH.3RD\_STAGE.TK\_OUTCOMES\_INCORRECTLY\_PREDICTED counts the number of incorrect not-taken predicate predictions on taken branches made by the BAC in the third stage of the core pipeline, including overrides of TAR taken predictions (made in the first stage) on the last instances of loop-closing branches
- **Event Code:** 0x0F, **Umask:** 1101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### BRANCH\_PREDICTOR.ALL.ALL\_PREDICTIONS

- **Title:** All Branch Predictions, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.ALL.ALL\_PREDICTIONS counts all branch predictions, which take place in the front-end of the processor. Note that this number does not necessarily equal the total number of branch instructions in the code, as branch predictions are made on a bundle basis (i.e., there is only one prediction per multiway branch bundle)
- **Event Code:** 0x10, **Umask:** 0000, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### BRANCH\_PREDICTOR.ALL.CORRECT\_PREDICTIONS

- **Title:** Correct Branch Predictions by All Predictors, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.ALL.CORRECT\_PREDICTIONS counts all branch predictions that do not necessitate a back-end branch misprediction flush, independent of predictor. A mismatch between the predicted and actual values of the branch predicate or target results in a branch misprediction. Return branches must additionally predict privilege level and previous function state
- **Event Code:** 0x10, **Umask:** 0001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### BRANCH\_PREDICTOR.ALL.WRONG\_PATH

- **Title:** Incorrect Predicate Predictions by All Predictors, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.ALL.WRONG\_PATH counts branch mispredictions that result from a mismatch of the predicted and actual values of the branch predicate, independent of predictor
- **Event Code:** 0x10, **Umask:** 0010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### BRANCH\_PREDICTOR.ALL.WRONG\_TARGET

- **Title:** Incorrect Target Predictions by All Predictors, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.ALL.WRONG\_TARGET counts branch mispredictions that result from a mismatch of the predicted and actual values of the branch target, independent of predictor
- **Event Code:** 0x10, **Umask:** 0011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PREDICTOR.1ST\_STAGE.ALL\_PREDICTIONS**

- **Title:** All Branch Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.1ST\_STAGE.ALL\_PREDICTIONS counts the number of branch predictions made in the first stage of the core pipeline by the TAR. The TAR is the only predictor operating in that stage of the pipeline and it only makes taken predictions. The PLP in the third stage may override a TAR predicate prediction on a loop-closing branch. The prediction flow is as follows:

```

if (TAR Hit)
    monitor++
    Read Target from TAR

```

- **Event Code:** 0x10, **Umask:** 0100, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PREDICTOR.1ST\_STAGE.CORRECT\_PREDICTIONS**

- **Title:** Correct Branch Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.1ST\_STAGE.CORRECT\_PREDICTIONS counts the number of branches correctly predicted taken by the TAR, both in predicate and target
- **Event Code:** 0x10, **Umask:** 0101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PREDICTOR.1ST\_STAGE.WRONG\_PATH**

- **Title:** Incorrect Predicate Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.1ST\_STAGE.WRONG\_PATH counts the number of actually not-taken branches predicted by the TAR (excluding overrides by the PLP)
- **Event Code:** 0x10, **Umask:** 0110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

**BRANCH\_PREDICTOR.1ST\_STAGE.WRONG\_TARGET**

- **Title:** Incorrect Target Predictions made in the first pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.1ST\_STAGE.WRONG\_TARGET counts the number of taken branches that were resteeered to an incorrect target by the TAR
- **Event Code:** 0x10, **Umask:** 0111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

## BRANCH\_PREDICTOR.2ND\_STAGE.ALL\_PREDICTIONS

- **Title:** All Branch Predictions in the second pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.2ND\_STAGE.ALL\_PREDICTIONS counts the number of branch predictions made in the second stage of the core pipeline. The following structures operate in that stage: BPT and MBPT (for predicates), TAC and RSB (for targets). Predictions are made in the second stage only if no predictions were made during the first stage. Any prediction made in this stage will be counted, except when a taken predicate prediction is made by the BPT/MBPT on a non-return branch and no target is available from the TAC. The branch prediction structures interact in the following manner:

```

if ((BPT Hit) or (MBPT Hit))
  if (Predicted Taken)
    if (Predicted Return Branch)
      monitor++
      Read Target from RSB
    else
      if (TAC Hit)
        monitor++
        Read Target from TAC
      else
        Get Target from BAC in the 3rd Stage
  else
    monitor++
    Follow Sequential Path
else
  if (TAC Hit)
    monitor++
    Read Target from TAC
  else
    Follow Sequential Path

```

- **Event Code:** 0x10, **Umask:** 1000, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

## BRANCH\_PREDICTOR.2ND\_STAGE.CORRECT\_PREDICTIONS

- **Title:** Correct Branch Predictions made in the second pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.2ND\_STAGE.CORRECT\_PREDICTIONS counts the number of correct predicate predictions made by the BPT/MBPT or the TAC in the second stage of the core pipeline. If the predicate prediction is taken, the correct target must be provided during that stage by the RSB or the TAC. Correct taken predicate predictions made by the BPT/MBPT on non-return branches that miss the TAC require the BAC to provide a target in the third stage and are not counted by this monitor
- **Event Code:** 0x10, **Umask:** 1001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### BRANCH\_PREDICTOR.2ND\_STAGE.WRONG\_PATH

- **Title:** Incorrect Predicate Predictions made in the second pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.2ND\_STAGE.WRONG\_PATH counts the number of incorrect not-taken predicate predictions made in the second stage of the core pipeline, and the number of incorrect taken predicate predictions made in that stage if a target was also provided
- **Event Code:** 0x10, **Umask:** 1010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### BRANCH\_PREDICTOR.2ND\_STAGE.WRONG\_TARGET

- **Title:** Incorrect Target Predictions made in the second pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.2ND\_STAGE.WRONG\_TARGET counts the number of branches that were correctly predicted taken by the BPT/MBPT or TAC, but were resteeered to an incorrect target by the RSB or the TAC
- **Event Code:** 0x10, **Umask:** 1011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### BRANCH\_PREDICTOR.3RD\_STAGE.ALL\_PREDICTIONS

- **Title:** All Branch Predictions made in the third pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.3RD\_STAGE.ALL\_PREDICTIONS counts the number of branch predictions made in the third stage of the core pipeline by the BAC. The BAC can make both predicate predictions (based on the whether hint field of the branch) and target predictions, in the following manner:

```

if (TAR Hit)
    if (Predicted Last Instance of Loop-Closing Branch)
        monitor++
        PLP Override of TAR Taken Prediction
        Resteer Frontend to Sequential Address
else
    if ((BPT Hit) or (MBPT Hit))
        if (Predicted Taken)
            if (not (TAC Hit))
                if (not (Predicted Return Branch))
                    monitor++
                    Compute Target
        else
            if (not (TAC Hit))
                monitor++
                Read Whether Hint Field for Predicate Prediction
                if (Predicted Taken)
                    Read BType Field for Type Information
                    if (Indirect Branch)
                        Read Target from RSB
                    else
                        Compute Target
            else
                Follow Sequential Path

```

- **Event Code:** 0x10, **Umask:** 1100, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PREDICTOR.3RD\_STAGE.CORRECT\_PREDICTIONS**

- **Title:** Correct Branch Predictions made in the third pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.3RD\_STAGE.CORRECT\_PREDICTIONS counts the number of correct branch predictions made by the BAC, including target predictions of branches whose predicate was supplied by a different predictor. For predicted-taken branches, both predicate and target must be correct
- **Event Code:** 0x10, **Umask:** 1101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PREDICTOR.3RD\_STAGE.WRONG\_PATH**

- **Title:** Incorrect Predicate Predictions made in the third pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.3RD\_STAGE.WRONG\_PATH counts branches whose predicate was incorrectly predicted by the BAC (based on the whether hint field of the branch), and not-taken branches whose taken predicate prediction by another predictor caused the BAC to supply a target
- **Event Code:** 0x10, **Umask:** 1110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

### **BRANCH\_PREDICTOR.3RD\_STAGE.WRONG\_TARGET**

- **Title:** Incorrect Target Predictions made in the third pipeline stage, **Category:** Branch
- **Definition:** BRANCH\_PREDICTOR.3RD\_STAGE.WRONG\_TARGET counts taken branches that were correctly predicted taken by any predictor, but whose target was incorrectly supplied by the BAC
- **Event Code:** 0x10, **Umask:** 1111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

## BRANCH\_TAKEN\_SLOT

- **Title:** Taken Branch Detail, **Category:** Branch
- **Definition:** BRANCH\_TAKEN\_SLOT monitors which slot number in a branch bundle (single-way or multiway) a taken branch occupies, or records that there were no taken branches in the given branch bundle. Use this monitor behind the downstream opcode matcher, rather than IA64\_TAGGED\_INST\_RETIRED, to count dynamic br.calls and br.rets.
- **Event Code:** 0x0D, **Umask:** See below, **PMC/PMD:** 4,5,6,7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

The SLOT\_MASK unit mask defined by [Table 7-29](#) allows profiling of taken branches based on their instruction slot number. If multiple bits are set in the SLOT\_MASK, all the set cases are included in the event count. The processor uses the following equation to determine the event outcome in each cycle:

```

(PMC.umask{16}
  and (branch in slot 0 is taken))
or (PMC.umask{17}
  and (branch in slot 0 is NOT taken)
  and (branch in slot 1 is taken))
or (PMC.umask{18}
  and (branch in slot 0 is NOT taken)
  and (branch in slot 1 is NOT taken)
  and (branch in slot 2 is taken))
or (PMC.umask{19}
  and (branch in slot 0 is NOT taken)
  and (branch in slot 1 is NOT taken)
  and (branch in slot 2 is NOT taken))

```

**Table 7-29. Slot unit mask for BRANCH\_TAKEN\_SLOT**

SLOT_MASK	PMC.umask {19:16}	Description
Instruction Slot 0	xxx1	Count if branch in slot 0 is first taken branch
Instruction Slot 1	xx1x	Count if branch in slot 1 is first taken branch
Instruction Slot 2	x1xx	Count if branch in slot 2 is first taken branch
No taken branch	1xxx	Count if NO branch was taken

## BUS\_ALL

- **Title:** Bus Transactions **Category:** Frontside Bus
- **Definition:** BUS\_ALL counts all transactions issued on the bus. These include BRL, BRIL, BIL, BWL, BRP, BWP, IORD, IOWR, and the other transactions.
- **Event Code:** 0x47, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no



## BUS\_BRQ\_LIVE\_REQ\_LO and BUS\_BRQ\_LIVE\_REQ\_HI

- **Title:** BRQ Live Requests, **Category:** Frontside Bus
- **Definition:** BUS\_BRQ\_LIVE\_REQ counts the number of live entries in the bus request queue (BRQ). These events include L3 cache reads, BRL, BRIL, BRP, and IORD memory transactions. The count excludes cache line write backs, partial writes (BWP and IOWR) and write coalescing read for ownership transactions, since these have their own write queue. This performance monitor increments its count each core clock (not bus clock).
- **Event Code:** 0x5c (HI), 0x5b (LO), **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 2 (each)
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no  
 Two hardware events are needed to count this:  
 BUS\_BRQ\_LIVE\_REQ\_HI - counts the two most significant bit of the 4-bit outstanding BRQ request count  
 BUS\_BRQ\_LIVE\_REQ\_LO - counts the two least significant bit of the 4-bit outstanding BRQ request count

## BUS\_BRQ\_READ\_REQ\_INSERTED

- **Title:** BRQ Requests Inserted **Category:** Frontside Bus
- **Definition:** BUS\_BRQ\_READ\_REQ\_INSERTED counts the number of reads (BRL) and read for ownership (BRIL) requests that are inserted into the BRQ. The count excludes cache line write backs, partial and coalescing writes, since these have their own write queue.
- **Event Code:** 0x5d, **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no  
 BRQ inserts can be used to directly measure combined far cache/BUS latencies as follows:  

$$\text{avg\_brq\_req\_outstanding\_per\_cycle} = (\text{BUS\_BRQ\_LIVE\_REQ} / \text{delta}(\text{cycles}))$$

$$\text{avg\_brq\_latency} = (\text{BUS\_BRQ\_LIVE\_REQ} / \text{BUS\_BRQ\_READ\_REQ\_INSERTED})$$
 The only caveat is that the tracked BRQ inserts holds read and read for ownership, but not write coalescing write backs.

**BUS\_BURST**

- **Title:** Bus Burst Transactions **Category:** Frontside Bus
- **Definition:** BUS\_BURST counts the number of full cache line (burst mode) bus memory transactions. These include BRL, BRIL, and BWL transactions.
- **Event Code:** 0x49, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_HITM**

- **Title:** Bus Hit Modified Line Transactions **Category:** Frontside Bus
- **Definition:** BUS\_HITM counts the number of memory transactions which caused HITM to be asserted. The following memory transactions are included in the performance monitor: BRL, BWL, BRIL, and BIL. Only events originated by this processor are counted.
- **Event Code:** 0x44, **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_IO**

- **Title:** IA-32 Compatible I/O Bus Transactions **Category:** Frontside Bus
- **Definition:** BUS\_IO counts the number of I/O transactions. These include either IORD or IOWR transactions.
- **Event Code:** 0x50, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**Note:** only the “Any” and “Self” unit masks are supported.

## BUS\_IOQ\_LIVE\_REQ\_LO and BUS\_IOQ\_LIVE\_REQ\_HI

- **Title:** In-Order Bus Queue Results, **Category:** Frontside Bus
- **Definition:** BUS\_IOQ\_LIVE\_REQ counts the number of live bus requests in the in order bus queue. This performance monitor increments its count each core clock (not bus clock).
- **Event Code:** 0x58 (HI), 0x57 (LO), **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 3 (each)
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no  
 BUS\_IOQ\_LIVE\_REQ can then be computed as  $BUS\_IOQ\_LIVE\_REQ\_HI * 4 + BUS\_IOQ\_LIVE\_REQ\_LO$   
 Two hardware events are needed to count this:  
 BUS\_IOQ\_LIVE\_REQ\_HI - counts the two most significant bit write backs of the 4-bit outstanding IOQ request count  
 BUS\_IOQ\_LIVE\_REQ\_LO - counts the two least significant bits of the 4-bit outstanding IOQ request count

## BUS\_LOCK

- **Title:** IA-32 Compatible Bus Lock Transactions **Category:** Frontside Bus
  - **Definition:** BUS\_LOCK counts the number of IA-32 compatible bus lock transactions.
  - **Event Code:** 0x53, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
  - **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no
- Note:** Only the “Any” unit mask is supported.

## BUS\_LOCK\_CYCLES

- **Title:** IA-32 Compatible Bus Lock Cycles **Category:** Frontside Bus
  - **Definition:** BUS\_LOCK\_CYCLES counts the number of bus clocks that the bus is locked due to IA-32 compatible bus lock transactions.
  - **Event Code:** 0x54, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
  - **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no
- Note:** Only the “Any” unit mask is supported.

**BUS\_MEMORY**

- **Title:** Bus Memory Transactions **Category:** Frontside Bus
- **Definition:** BUS\_MEMORY counts the number of bus memory transactions. These include BRL, BRIL, BIL, BWL, BRP, and BWP transactions.
- **Event Code:** 0x4a, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_PARTIAL**

- **Title:** Bus Partial Transactions **Category:** Frontside Bus
- **Definition:** BUS\_PARTIAL counts the number of partial bus memory transactions. These include BRP and BWP transactions.
- **Event Code:** 0x48, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_RD\_ALL**

- **Title:** Bus Read Transactions **Category:** Frontside Bus
- **Definition:** BUS\_RD\_ALL counts the number of BRL memory transactions. These include both code and data BRL transactions.
- **Event Code:** 0x4b, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_RD\_DATA**

- **Title:** Bus Read Data Transactions **Category:** Frontside Bus
- **Definition:** BUS\_RD\_DATA counts the number of BRL data transactions.
- **Event Code:** 0x4c, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

### BUS\_RD\_HIT

- **Title:** Bus Read Hit Clean Non-local Cache Transactions **Category:** Frontside Bus
- **Definition:** BUS\_RD\_HIT counts the number of BRL memory transactions which caused HIT to be asserted. Only events originated by this processor are counted.
- **Event Code:** 0x40, **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

### BUS\_RD\_HITM

- **Title:** Bus Read Hit Modified Non-local Cache Transactions **Category:** Frontside Bus
- **Definition:** BUS\_RD\_HITM counts the number of BRL memory transactions which caused HITM to be asserted. Only events originated by this processor are counted.
- **Event Code:** 0x41, **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

### BUS\_RD\_INVALID

- **Title:** Bus Read Invalidated Line **Category:** Frontside Bus
- **Definition:** BUS\_RD\_INVALID counts the number of BIL memory transactions. On Itanium processors, these transactions are only generated from flush cache instructions.
- **Event Code:** 0x4e, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

### BUS\_RD\_INVALID\_BST

- **Title:** Bus BRIL Burst Transactions **Category:** Frontside Bus
- **Definition:** BUS\_RD\_INVALID counts the number of BRIL memory transactions. These transactions are typically generated from memory stores, RFO (read for ownership) events.
- **Event Code:** 0x4f, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_RD\_INVALID\_HITM**

- **Title:** Bus BIL Transaction Results in HITM **Category:** Frontside Bus
- **Definition:** BUS\_RD\_INVALID\_HITM counts the number of BIL transactions which caused HITM to be asserted. Only events originated by this processor are counted.
- **Event Code:** 0x42, **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_RD\_INVALID\_BST\_HITM**

- **Title:** Bus BRIL Burst Transaction Results in HITM **Category:** Frontside Bus
- **Definition:** BUS\_RD\_INVALID\_BST\_HITM counts the number of BRIL transactions which caused HITM to be asserted. Only events originated by this processor are counted.
- **Event Code:** 0x43, **Umask:** Ignored, **PMC/PMD:**4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_RD\_IO**

- **Title:** IA-32 Compatible I/O Read Transactions **Category:** Frontside Bus
- **Definition:** BUS\_RD\_IO counts the number of IORD transactions.
- **Event Code:** 0x51, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**Note:** Only the “Any” and “Self” unit masks are supported.

**BUS\_RD\_PRTL**

- **Title:** Bus Read Partial Transactions **Category:** Frontside Bus
- **Definition:** BUS\_RD\_PRTL counts the number of partial read memory transactions. These include BRP transactions.
- **Event Code:** 0x4d, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

## BUS\_SNOOPS

- **Title:** Bus Snoops Total **Category:** Frontside Bus
- **Definition:** BUS\_SNOOPS counts the number of internal snoops generated from bus memory transactions.
- **Event Code:** 0x46, **Umask:** See [Section 7.6.5, PMC/PMD:4, 5, 6, 7](#)  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**Note:** Only the “Any” unit mask is supported. “Any” counts the number of internal snoops generated from all bus transactions.

## BUS\_SNOOPS\_HITM

- **Title:** Bus Snoops Hit Modified Cache Line **Category:** Frontside Bus
- **Definition:** BUS\_SNOOPS\_HITM counts the number of internal snoops (generated from bus memory transactions) which hit a modified line in the local processor.
- **Event Code:** 0x45, **Umask:** See [Section 7.6.5, PMC/PMD:4, 5, 6, 7](#)  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**Note:** Only the “Any” mask is supported. “Any” counts the number of internal snoops, generated from all transactions, which hit a modified line.

## BUS\_SNOOP\_STALL\_CYCLES

- **Title:** Bus Snoop Stall Cycles **Category:** Frontside Bus
- **Definition:** BUS\_SNOOP\_STALL\_CYCLES counts the number of bus clocks that the bus is stalled due to snoop stalls.
- **Event Code:** 0x55, **Umask:** See [Section 7.6.5, PMC/PMD:4, 5, 6, 7](#)  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**Note:** Only “Self” and “Any” unit masks are supported. “Self” counts the number of snoop stalls generated due to memory transactions initiated by the local processor. “Any” counts all snoop stalls (those generated due to memory transactions initiated by the local processor, other processors, and the priority agent).

**BUS\_SNOOPQ\_REQ**

- **Title:** Bus Snoop Queue Requests, **Category:** Frontside Bus
- **Definition:** BUS\_SNOOPQ\_REQ counts the number of outstanding memory transactions that have not completed the snoop phase. This performance monitor increments its count each core clock (not bus clock). BUS\_SNOOPQ\_REQ is not equivalent to the number of valid entries in the snoop queue. This is due to the fact that entries can stay in the snoop queue beyond the snoop phase (e.g. for implicit write backs).
- **Event Code:** 0x56, **Umask:** Ignored, **PMC/PMD:**4, 5, **Max. Increment/Cycle:** 3
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**BUS\_WR\_WB**

- **Title:** Bus Write Back Transactions **Category:** Frontside Bus
- **Definition:** BUS\_WR\_WB counts the number of BWL memory transactions. These transactions are generated from either explicit write backs or coalescing writes. Currently, these will count BWL (if snoops are disabled).
- **Event Code:** 0x52, **Umask:** See [Section 7.6.5](#), **PMC/PMD:**4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**CPU\_CPL\_CHANGES**

- **Title:** Privilege level changes, **Category:** System
- **Definition:** CPU\_CPL\_CHANGES counts the number of privilege level changes
- **Event Code:** 0x34, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**CPU\_CYCLES**

- **Title:** CPU Cycles, **Category:** System
- **Definition:** CPU\_CYCLES counts elapsed processor cycles
- **Event Code:** 0x12, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no



## DATA\_ACCESS\_CYCLE

- **Title:** Data Access Stall Cycles, **Category:** Stall
- **Definition:** DATA\_ACCESS\_CYCLE counts the number of cycles that the pipeline is stalled or flushed due to instructions waiting for data on cache misses, L1D way mispredictions, and DTC misses.
- **Event Code:** 0x03, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

## DATA\_EAR\_EVENTS

- **Title:** L1 Data Cache EAR Events, **Category:** L1 Data Cache
- **Definition:** DATA\_EAR\_EVENTS counts the number of data cache or DTLB events captured by the Data Cache Unit Event Address Register
- **Event Code:** 0x67, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

## DATA\_REFERENCES\_RETIRED

- **Title:** Retired Data Memory References, **Category:** L1 Data Cache
- **Definition:** DATA\_REFERENCES\_RETIRED counts the number of data memory references retired by the processor memory pipeline. The count includes check loads, uncacheable accesses, RSE operations, VHPT memory references, semaphores, and FP memory references. Predicated off operations are excluded
- **Event Code:** 0x63, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: yes

## DEPENDENCY\_ALL\_CYCLE

- **Title:** Scoreboard Dependency and Dispersal Break Cycles, **Category:** Stall
- **Definition:** DEPENDENCY\_ALL\_CYCLE counts the number of cycles attributable to data (scoreboard) dependency on integer or FP operations (not counting cache/memory access), or issue-limit stalls (e.g., implicit and explicit stops). Floating-point flushes and delays due to control and application register reads and writes are factored in as well.
- **Event Code:** 0x06, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**DEPENDENCY\_SCOREBOARD\_CYCLE**

- **Title:** Scoreboard Dependency Cycles, **Category:** Stall
- **Definition:** DEPENDENCY\_SCOREBOARD\_CYCLE counts the number of cycles attributable to data (scoreboard) dependency on integer or FP operations (not counting cache/memory access). Floating-point flushes and delays due to control and application register reads and writes are factored in as well.
- **Event Code:** 0x02, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**DTC\_MISSES**

- **Title:** DTC Misses, **Category:** System
- **Definition:** DTC\_MISSES counts the number of DTC misses for data requests
- **Event Code:** 0x60, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

**DTLB\_INSERTS\_HPW**

- **Title:** Hardware Page Walker Inserts into the DTLB, **Category:** System
- **Definition:** DTLB\_INSERTS\_HPW counts the number of DTLB inserts completed by the hardware page table walker
- **Event Code:** 0x62, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

**DTLB\_MISSES**

- **Title:** DTLB Misses, **Category:** System
- **Definition:** DTLB\_MISSES counts the number of DTLB misses for demand requests
- **Event Code:** 0x61, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

## EXPL\_STOPS\_DISPERSED

- **Title:** Explicit Stops Dispersed, **Category:** Instruction Issue
- **Definition:** EXPL\_STOPS\_DISPERSED counts the number of explicit programmer-specified stops, including those encountered during hardware speculative wrong-path execution
- **Event Code:** 0x2E, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

## FP\_OPS\_RETIRE\_HI

- **Title:** FP Operations Retired (High), **Category:** Execution
- **Definition:** FP\_OPS\_RETIRE\_HI and FP\_OPS\_RETIRE\_LO together compute the derived event FP\_OPS\_RETIRE.d which is the weighted sum of retired FP operations
- **Event Code:** 0x0A, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 3
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

FP\_OPS\_RETIRE.d, a derived value, is computed as  $FP\_OPS\_RETIRE\_HI * 4 + FP\_OPS\_RETIRE\_LO$ . Weights for individual FP ops:  $f_{norm}=1$ ,  $f_{add}=1$ ,  $f_{mpy}=1$ ,  $f_{ma}=2$ ,  $f_{ms}=2$ ,  $f_{sub}=1$ ,  $f_{pma}=4$ ,  $f_{pmpy}=4$ ,  $f_{pms}=4$ ,  $f_{nma}=2$ ,  $f_{rcpa}=1$ ,  $f_{rsqrta}=1$ ,  $f_{pnma}=4$ ,  $f_{prcpa}=2$ ,  $f_{prsqrta}=2$ ,  $x_{ma}=0$

**Note:** Integer multiply instructions ( $x_{ma}$ ) are not counted as floating-point operations (even though they are executed in the floating-point multiplier).

## FP\_OPS\_RETIRE\_LO

- **Title:** FP Operations Retired (Low), **Category:** Execution
- **Definition:** FP\_OPS\_RETIRE\_HI and FP\_OPS\_RETIRE\_LO together compute the derived event FP\_OPS\_RETIRE.d which is the weighted sum of retired FP operations
- **Event Code:** 0x09, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 3
- **Qualification:** See FP\_OPS\_RETIRE\_HI on page 91

## FP\_FLUSH\_TO\_ZERO

- **Title:** FP Result Flushed to Zero, **Category:** Execution
- **Definition:** FP\_FLUSH\_TO\_ZERO counts the number of times a near zero result is flushed to zero in FTZ mode. Parallel FP operations which cause one or both results to flush to zero will increment the event count only by one (i.e. even if both results are flushed to zero)
- **Event Code:** 0x0B, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

**FP\_SIR\_FLUSH**

- **Title:** FP SIR Flushes, **Category:** Execution
- **Definition:** FP\_SIR\_FLUSH counts the number of times a Safe Instruction Recognition (SIR) flush occurs
- **Event Code:** 0x0C, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

**IA32\_INST\_RETIRED**

- **Title:** Retired IA-32 Instructions, **Category:** System
- **Definition:** IA32\_INST\_RETIRED counts the number of IA-32 instructions retired
- **Event Code:** 0x15, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**IA64\_INST\_RETIRED**

- **Title:** Retired Itanium Instructions, **Category:** Execution
- **Definition:** IA64\_INST\_RETIRED counts all retired Itanium instructions. The count includes predicated on and off instructions, NOPs, but excludes hardware-inserted RSE operations. This event is equal to IA64\_TAGGED\_INST\_RETIRED with a zero unit mask
- **Event Code:** 0x08, **Umask:** 0000, **PMC/PMD:** 4, 5 **Max. Increment/Cycle:** 6
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

**IA64\_TAGGED\_INST\_RETIRED**

- **Title:** Retired Tagged Itanium Instructions **Category:** Execution
- **Definition:** IA64\_TAGGED\_INST\_RETIRED is analogous to IA64\_INST\_RETIRED, except that it further qualifies event selection with the instruction address range and opcode match settings in the IBR and PMC registers
- **Event Code:** 0x08 **Umask:** See below **PMC/PMD:** 4, 5 **Max. Increment/Cycle:** 6
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

The TAG\_SELECT unit mask defined in [Table 7-30](#) always qualifies the event count of IA64\_TAGGED\_INST\_RETIRED with either the opcode match register PMC<sub>8</sub> or PMC<sub>9</sub>. Note that the setting of PMC<sub>8</sub> qualifies all down-stream event monitors. To ensure that other monitored events are counted independent of the opcode matcher, all mifb and all mask bits of PMC<sub>8</sub> should be set to one (all opcodes match). The settings of PMC<sub>9</sub> do not affect other event monitors.

Also, note that umask 0011 is distinct in that it also counts, in addition to instructions matched by the appropriate opcode matcher, architecturally invisible RSE fills and spills when the parent instruction (such as an alloc or br.ret) causing them is matched by the combination in PMC8. Thus, the difference in counts obtained between using PMC8 and PMC9 as opcode matchers is the amount of RSE activity.

**Table 7-30. Retired Event Selection by Opcode Match**

TAG_SELECT	PMC.umask {19:16}	Description
PMC <sub>8</sub> tag	0011	Instruction tagged by opcode matcher PMC <sub>8</sub>
PMC <sub>9</sub> tag	0010	Instruction tagged by opcode matcher PMC <sub>9</sub>
All	0000	All retired instructions (regardless of whether they were tagged or not)
Undefined	All other umask settings	Undefined event count

### INST\_ACCESS\_CYCLE

- **Title:** Instruction Access Cycles, **Category:** Stall
- **Definition:** INST\_ACCESS\_CYCLE counts the number of cycles where there are no back-end stalls or flushes, the decoupling buffer is empty, and the front-end is stalled waiting on an L1I or ITLB miss.
- **Event Code:** 0x01, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### INST\_DISPERSED

- **Title:** Instructions Dispersed, **Category:** Instruction Issue
- **Definition:** INST\_DISPERSED counts the number of instructions dispersed (including nops) from the front-end to the back-end of the machine. The count includes instruction dispersal on the wrong execution path; i.e., in the shadow of a branch misprediction flush or other back-end flush
- **Event Code:** 0x2D, **Umask:** Ignored, **PMC/PMD:** 4, 5 **Max. Increment/Cycle:** 6
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

### INST\_FAILED\_CHKS\_RETIRED.ALL

- **Title:** Failed Speculative Check Loads, **Category:** Execution
- **Definition:** INST\_FAILED\_CHKS\_RETIRED.ALL counts the number of failed speculative check load instructions (chk.s). The count excludes predicated off chk.s instructions and includes both integer and FP variants
- **Event Code:** 0x35, **Umask:** xx11, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

**INST\_FAILED\_CHKS\_RETIRED.FP**

- **Title:** Failed Speculative FP Check Loads, **Category:** Execution
- **Definition:** INST\_FAILED\_CHKS\_RETIRED.FP counts the number of failed speculative check load instructions (`chk . s`). The count excludes predicated off `chk . s` instructions and includes only FP variants
- **Event Code:** 0x35, **Umask:** xx10, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no

**INST\_FAILED\_CHKS\_RETIRED.INTEGER**

- **Title:** Failed Speculative Integer Check Loads, **Category:** Execution
- **Definition:** INST\_FAILED\_CHKS\_RETIRED.INTEGER counts the number of failed speculative check load instructions (`chk . s`). The count excludes predicated off `chk . s` instructions and includes only integer variants
- **Event Code:** 0x35, **Umask:** xx01, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: no
- **Qualification:**

**INSTRUCTION\_EAR\_EVENTS**

- **Title:** Instruction EAR Events, **Category:** Instruction Cache
- **Definition:** INSTRUCTION\_EAR\_EVENTS counts the number of EAR captures for L1I and ITLB events
- **Event Code:** 0x23, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

**ISA\_TRANSITIONS**

- **Title:** Itanium ISA to IA-32 ISA Transitions, **Category:** System
- **Definition:** ISA\_TRANSITIONS counts the number of instruction set transitions from Itanium ISA to IA-32. This is the number of times the PSR.is bit toggles from 0 to 1 due to `br . ia` or `r fi` to IA-32 code
- **Event Code:** 0x14, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### ISB\_LINES\_IN

- **Title:** Instruction Streaming Buffer Lines In, **Category:** Instruction Cache
- **Definition:** ISB\_LINES\_IN counts the number of 32-byte L1I cache lines written from L2 (and beyond) into the Instruction Streaming Buffer as a consequence of instruction demand miss and instruction prefetch requests
- **Event Code:** 0x26, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### ITLB\_INSERTS\_HPW

- **Title:** Hardware Page Walker Inserts into the ITLB, **Category:** System
- **Definition:** ITLB\_INSERTS\_HPW counts the number of ITLB inserts done by the hardware page table walker
- **Event Code:** 0x28, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

### ITLB\_MISSES\_FETCH

- **Title:** ITLB Demand Misses, **Category:** System
- **Definition:** ITLB\_MISSES\_FETCH counts the number of demand ITLB misses
- **Event Code:** 0x27, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

### L1D\_READ\_FORCED\_MISSES\_RETIRED

- **Title:** L1 Data Cache Forced Load Misses, **Category:** L1 Data Cache
- **Definition:** L1D\_READ\_FORCED\_MISSES\_RETIRED counts the number of loads that were forced to miss the L1 data cache due to memory ordering constraints, predicted L1 data cache misses, Store Buffer hits, or simultaneous L2 data returns to the register file
- **Event Code:** 0x6B, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: yes

**L1D\_READ\_MISSES\_RETIRE**

- **Title:** L1 Data Cache Read Misses, **Category:** L1 Data Cache
- **Definition:** L1D\_READ\_MISSES\_RETIRE counts the number of committed L1 data cache read misses. The count includes any read reference that could have been serviced by the L1 data cache (see L1D\_READS\_RETIRE event for a detailed list) but missed the cache. False misses are included in the event count. Since the L1 data cache is write-through, write misses are NOT counted
- **Event Code:** 0x66, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: yes

**L1D\_READS\_RETIRE**

- **Title:** L1 Data Cache Reads, **Category:** L1 Data Cache
- **Definition:** L1D\_READS\_RETIRE counts the number of committed L1 data cache reads (integer and RSE references). Excluded from the count are VHPT loads, check loads, L1 hinted loads, semaphores, uncacheable and FP loads. Predicated-off loads are also excluded, but wrong-path operations are included in the count
- **Event Code:** 0x64, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: yes

**L1I\_DEMAND\_READS**

- **Title:** L1I and ISB Instruction Demand Lookups, **Category:** Instruction Cache
- **Definition:** L1I\_DEMAND\_READS counts the number of 32-byte instruction demand L1I/ISB lookups, independent of the hit/miss outcome
- **Event Code:** 0x20, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no  
Qualifications based on instruction address range may be inaccurate

**L1I\_FILLS**

- **Title:** L1 Instruction Cache Fills, **Category:** Instruction Cache
- **Definition:** L1I\_FILLS counts the number of 32-byte lines moved from the Instruction Streaming Buffer into the L1 instruction cache
- **Event Code:** 0x21, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no



### L1I\_PREFETCH\_READS

- **Title:** L1I and ISB Instruction Prefetch Lookups, **Category:** Instruction Cache
- **Definition:** L1I\_PREFETCH\_READS counts the number of 64-byte instruction prefetch L1I/ISB lookups, independent of the hit/miss outcome.
- **Event Code:** 0x24, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

### L2\_DATA\_REFERENCES.ALL

- **Title:** L2 Data Read and Write References, **Category:** L2 Cache
- **Definition:** L2\_DATA\_REFERENCES.ALL counts all L2 data read and write accesses. The reported count is the number of requests prior to cache line merging. Semaphore operations are counted as one read and one write
- **Event Code:** 0x69, **Umask:** xx11, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: yes

### L2\_DATA\_REFERENCES.READS

- **Title:** L2 Data Read References, **Category:** L2 Cache
- **Definition:** L2\_DATA\_REFERENCES.READS counts all L2 data read accesses. The reported count is the number of requests prior to cache line merging. Semaphore operations are counted as one read
- **Event Code:** 0x69, **Umask:** xx01, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: yes

### L2\_DATA\_REFERENCES.WRITES

- **Title:** L2 Data Write References, **Category:** L2 Cache
- **Definition:** L2\_DATA\_REFERENCES.WRITES counts all L2 data write accesses. The reported count is the number of requests prior to cache line merging. Semaphore operations are counted as one write
- **Event Code:** 0x69, **Umask:** xx10, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: yes, Data Address Range: yes

## L2\_FLUSH\_DETAILS

- **Title:** L2 Flush Details, **Category:** L2 Cache
- **Definition:** L2\_FLUSH\_DETAILS allows a detailed breakdown of L2 pipeline flushes by cause. This event counts the number of L2 pipeline flushes constrained by the conditions specified in the 4-bit unit mask defined by [Table 7-31 on page 98](#). All combinations of the four unit mask bits are supported
- **Event Code:** 0x77, **Umask:** See below, **PMC/PMD:** 4, 5, 6, 7, **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**Table 7-31. Unit Mask Bits {19:16} for L2\_FLUSH\_DETAILS Event**

L2 Flush Reason	PMC.umask {19:16}	Description
L2_ST_BUFFER_FLUSH	xxx1	L2 store to store conflict due to (a) Same store buffer entry (b) Back to back stores
L2_ADDR_CONFLICT	xx1x	L2 flushed due to MESI update on load follows store
L2_BUS_REJECT	x1xx	L2 flushed due to bus constraints
L2_FULL_FLUSH	1xxx	L2 flushed due to one of: (a) Store buffer full (b) Load miss buffer full

## L2\_FLUSHES

- **Title:** L2 Flushes, **Category:** L2 Cache
- **Definition:** L2\_FLUSHES counts the number of L2 pipeline flushes due to Store Buffer conflicts, address conflicts, full L3 and bus queues, and other such reasons
- **Event Code:** 0x76, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

## L2\_INST\_DEMAND\_READS

- **Title:** L2 Instruction Demand Fetch Requests, **Category:** Instruction Cache
- **Definition:** L2\_INST\_DEMAND\_READS counts the number of L2 instruction requests due to L1I demand fetch misses. The monitor counts the number of demand fetch look-ups that miss in both the L1I and the ISB, regardless of whether they hit or miss in the Request Address Buffer (RAB); i.e., the count includes misses to a line that has already been requested (secondary misses)
- **Event Code:** 0x22, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

## L2\_INST\_PREFETCH\_READS

- **Title:** L2 Instruction Prefetch Requests, **Category:** Instruction Cache
- **Definition:** L2\_INST\_PREFETCH\_READS counts all instruction prefetch requests issued to the unified L2 cache
- **Event Code:** 0x25, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

## L2\_MISSES

- **Title:** L2 Misses, **Category:** L2 Cache
- **Definition:** L2\_MISSES counts the number of L2 cache misses (requests to uncacheable pages are excluded). The count includes misses caused by instruction fetch and prefetch, and data read and write operations. Secondary misses to the same L2 cache line will be counted as individual misses
- **Event Code:** 0x6A, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

## L2\_REFERENCES

- **Title:** L2 References, **Category:** L2 Cache
- **Definition:** L2\_REFERENCES counts the number of L2 cache references (requests to uncacheable pages are excluded). The count includes references by instruction fetch and prefetch, and data reads and writes. The maximum per-cycle increment is three: one instruction fetch and two data references
- **Event Code:** 0x68, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 3
- **Qualification:** Instruction Address Range: yes, Opcode matching: no, Data Address Range: no

## L3\_LINES\_REPLACED

- **Title:** L3 Cache Lines Replaced, **Category:** L3 Cache
- **Definition:** L3\_LINES\_REPLACED counts the number of valid L3 lines that have been victimized
- **Event Code:** 0x7F, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_MISSES**

- **Title:** L3 Misses, **Category:** L3 Cache
- **Definition:** L3\_MISSES counts the number of L3 misses. The number includes misses caused by both instruction and data requests and L2 line writebacks
- **Event Code:** 0x7C, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_READS.ALL\_READS.ALL**

- **Title:** Instruction and Data L3 Reads, **Category:** L3 Cache
- **Definition:** L3\_READS.ALL\_READS.ALL counts the number of all L3 read accesses, independent of the stream source (instruction or data) and the hit/miss outcome
- **Event Code:** 0x7D, **Umask:** 1111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_READS.ALL\_READS.HIT**

- **Title:** Instruction and Data L3 Read Hits, **Category:** L3 Cache
- **Definition:** L3\_READS.ALL\_READS.HIT counts the number of all L3 read hits, independent of the stream source (instruction or data)
- **Event Code:** 0x7D, **Umask:** 1101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_READS.ALL\_READS.MISS**

- **Title:** Instruction and Data L3 Read Misses, **Category:** L3 Cache
- **Definition:** L3\_READS.ALL\_READS.MISS counts the number of all L3 read misses, independent of the stream source (instruction or data)
- **Event Code:** 0x7D, **Umask:** 1110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_READS.DATA\_READS.ALL**

- **Title:** Data L3 Reads, **Category:** L3 Cache
- **Definition:** L3\_READS.DATA\_READS.ALL counts the number of data L3 read accesses, independent of the hit/miss outcome
- **Event Code:** 0x7D, **Umask:** 1011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_READS.DATA\_READS.HIT**

- **Title:** Data L3 Read Hits, **Category:** L3 Cache
- **Definition:** L3\_READS.DATA\_READS.HIT counts the number of data L3 read hits
- **Event Code:** 0x7D, **Umask:** 1001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_READS.DATA\_READS.MISS**

- **Title:** Data L3 Read Misses, **Category:** L3 Cache
- **Definition:** L3\_READS.DATA\_READS.MISS counts the number of data L3 read misses
- **Event Code:** 0x7D, **Umask:** 1010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_READS.INST\_READS.ALL**

- **Title:** Instruction L3 Reads, **Category:** L3 Cache
- **Definition:** L3\_READS.INST\_READS.ALL counts the number of instruction L3 read accesses, independent of the hit/miss outcome
- **Event Code:** 0x7D, **Umask:** 0111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_READS.INST\_READS.HIT**

- **Title:** Instruction L3 Read Hits, **Category:** L3 Cache
- **Definition:** L3\_READS.INST\_READS.HIT counts the number of instruction L3 read hits
- **Event Code:** 0x7D, **Umask:** 0101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_READS.INST\_READS.MISS**

- **Title:** Instruction L3 Read Misses, **Category:** L3 Cache
- **Definition:** L3\_READS.INST\_READS.MISS counts the number of instruction L3 read misses
- **Event Code:** 0x7D, **Umask:** 0110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_REFERENCES**

- **Title:** L3 References, **Category:** L3 Cache
- **Definition:** L3\_REFERENCES counts the number of L3 cache references (requests to uncacheable pages are excluded). The count includes references by instruction fetch and prefetch, data reads and writes, and L2 cache line most significant bit writebacks.
- **Event Code:** 0x7B, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_WRITES.ALL\_WRITES.ALL**

- **Title:** L3 Writes, **Category:** L3 Cache
- **Definition:** L3\_WRITES.ALL\_WRITES.ALL counts the number of L3 write accesses independent of the hit/miss outcome. The count includes both data writes and L2 writeback accesses (including L3 read for ownership requests that satisfy stores)
- **Event Code:** 0x7E, **Umask:** 1111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_WRITES.ALL\_WRITES.HIT**

- **Title:** L3 Write Hits, **Category:** L3 Cache
- **Definition:** L3\_WRITES.ALL\_WRITES.HIT counts the number of L3 write hits. The count includes both data writes and L2 writeback accesses (including L3 read for ownership requests that satisfy stores)
- **Event Code:** 0x7E, **Umask:** 1101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_WRITES.ALL\_WRITES.MISS**

- **Title:** L3 Write Misses, **Category:** L3 Cache
- **Definition:** L3\_WRITES.ALL\_WRITES.MISS counts the number of L3 write misses. The count includes both data writes and L2 writeback accesses (including L3 read for ownership requests that satisfy stores)
- **Event Code:** 0x7E, **Umask:** 1110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_WRITES.L2\_WRITEBACK.ALL**

- **Title:** L3 Writebacks, **Category:** L3 Cache
- **Definition:** L3\_WRITES.L2\_WRITEBACK.ALL counts the number of L3 write accesses that result from L2 writebacks, independent of hit/miss outcome
- **Event Code:** 0x7E, **Umask:** 1011, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_WRITES.L2\_WRITEBACK.HIT**

- **Title:** L3 Writeback Hits, **Category:** L3 Cache
- **Definition:** L3\_WRITES.L2\_WRITEBACK.HIT counts the number of L3 write hits that result from L2 writebacks
- **Event Code:** 0x7E, **Umask:** 1001, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_WRITES.L2\_WRITEBACK.MISS**

- **Title:** L3 Writeback Misses, **Category:** L3 Cache
- **Definition:** L3\_WRITES.L2\_WRITEBACK.MISS counts the number of L3 write misses that result from L2 writebacks
- **Event Code:** 0x7E, **Umask:** 1010, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_WRITES.DATA\_WRITES.ALL**

- **Title:** L3 Data Writes, **Category:** L3 Cache
- **Definition:** L3\_WRITES.DATA\_WRITES.ALL counts the number of L3 data write accesses independent of the hit/miss outcome
- **Event Code:** 0x7E, **Umask:** 0111, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

### **L3\_WRITES.DATA\_WRITES.HIT**

- **Title:** L3 Data Write Hits, **Category:** L3 Cache
- **Definition:** L3\_WRITES.DATA\_WRITES.HIT counts the number of L3 data write hits
- **Event Code:** 0x7E, **Umask:** 0101, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**L3\_WRITES.DATA\_WRITES.MISS**

- **Title:** L3 Data Write Misses, **Category:** L3 Cache
- **Definition:** L3\_WRITES.DATA\_WRITES.MISS counts the number of L3 data write misses
- **Event Code:** 0x7E, **Umask:** 0110, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode matching: no, Data Address Range: no

**LOADS\_RETIRE**

- **Title:** Retired Loads, **Category:** Memory
- **Definition:** LOADS\_RETIRE counts the number of retired loads. The count includes integer, FP, RSE, VHPT, uncacheable loads and failed check loads (ld.c). Check loads that hit in the ALAT are *not* counted. Predicated-off operations are not counted
- **Event Code:** 0x6C, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

**MEMORY\_CYCLE**

- **Title:** Combined Memory Stall Cycles, **Category:** Stall
- **Definition:** MEMORY\_CYCLE counts the number of cycles that the pipeline is stalled or flushed due to instructions waiting for data on cache misses, L1D way mispredictions, DTC misses, and RSE traffic.
- **Event Code:** 0x07, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**MISALIGNED\_LOADS\_RETIRE**

- **Title:** Retired Unaligned Load Instructions, **Category:** Memory
- **Definition:** MISALIGNED\_LOADS\_RETIRE counts the number of retired unaligned loads that the hardware handled. The count includes integer, FP, and failed check loads (ld.c). Check loads that hit in the ALAT are *not* counted. Predicated-off operations are not counted
- **Event Code:** 0x70, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes



## MISALIGNED\_STORES\_RETIRED

- **Title:** Retired Unaligned Store Instructions, **Category:** Memory
- **Definition:** MISALIGNED\_STORES\_RETIRED counts the number of retired unaligned store instructions that the hardware handled. The count includes integer, FP, and uncacheable stores. Predicated-off operations are not counted
- **Event Code:** 0x71, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

## NOPS\_RETIRED

- **Title:** Retired Nop Instructions, **Category:** Execution
- **Definition:** NOPS\_RETIRED counts the number of retired nop.i, nop.m or nop.b instructions. The count excludes predicated off nop instructions
- **Event Code:** 0x30, **Umask:** Ignored, **PMC/PMD:** 4, 5 **Max. Increment/Cycle:** 6
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

## PIPELINE\_ALL\_FLUSH\_CYCLE

- **Title:** Combination of Pipeline Flush Cycles caused by either a front-end or a back-end source, **Category:** Stall
- **Definition:** PIPELINE\_ALL\_FLUSH\_CYCLE, for a given cycle, either counts the number of cycles spent during a front-end re-steer of the pipeline (due to a correctly predicted taken branch), or counts the number of cycles spent during certain back-end re-steers (due to a branch misprediction, ALAT flush or exception/serialization flush). This monitor does not count DTC flushes, way mispredictions, or floating-point flushes.
- **Event Code:** 0x04, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

## PIPELINE\_BACKEND\_FLUSH\_CYCLE

- **Title:** Combination of Pipeline Flush Cycles caused by either a Branch Misprediction or an Exception **Category:** Stall
- **Definition:** PIPELINE\_BACKEND\_FLUSH\_CYCLE counts the number of cycles spent during back-end re-steers of the pipeline (due to a branch misprediction, ALAT flush or exception/serialization flush). This monitor does not count DTC flushes, way mispredictions, or floating-point flushes.
- **Event Code:** 0x00, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

## PIPELINE\_FLUSH

- **Title:** Pipeline Flush, **Category:** System
- **Definition:** PIPELINE\_FLUSH counts how often the Itanium processor pipeline is flushed due to IEU bypass conflict (caused by non-unit latency MMX operations such as variable shifts), data translation cache miss, L1 data cache way mispredict or other reasons such as an exception flush or an instruction serialization. Combinations of different flush reasons may be chosen by appropriately setting the umask. The monitor does not include branch misprediction flushes
- **Event Code:** 0x33, **Umask:** See below, **PMC/PMD:** 4, 5, 6, 7  
**Max. Increment/Cycle:** 1
- **Qualification:** Instruction Address Range: no, Opcode Matching: no, Data Address Range: no

**Table 7-32. Unit Mask Bits {19:18} for PIPELINE\_FLUSH Event**

FLUSH_TYPE	PMC.umask {19:16}	Description
IEU_FLUSH	1xxx	IEU bypass flush
DTC_FLUSH	x1xx	Data Translation Cache Miss flush
L1D_WAYMP_FLUSH	xx1x	L1 Way Misprediction flush
OTHER_FLUSH	xxx1	Other flush reason: exception flush or an instruction serialization.

## PREDICATE\_SQUASHED\_RETIRE

- **Title:** Instructions Squashed Due to Predicate Off, **Category:** Execution
- **Definition:** PREDICATE\_SQUASHED\_RETIRE counts the number of instructions squashed due to a false qualifying predicate. The count includes all predicated off nops except nop.b's. Predicated off B-syllables (including nop.b) are not counted
- **Event Code:** 0x31, **Umask:** Ignored, **PMC/PMD:** 4, 5 **Max. Increment/Cycle:** 6
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: no

## RSE\_LOADS\_RETIRE

- **Title:** RSE Load Accesses, **Category:** Execution
- **Definition:** RSE\_LOADS\_RETIRE counts the number of retired RSE loads
- **Event Code:** 0x72, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Refer to RSE\_REFERENCES\_RETIRE on page 107

## RSE\_REFERENCES\_RETIRED

- **Title:** RSE Accesses, **Category:** Execution
- **Definition:** RSE\_REFERENCES\_RETIRED counts the number of retired RSE loads and stores
- **Event Code:** 0x65, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

RSE loads and stores are considered tagged if the `alloc`, `loadrs`, `flushrs` or branch return or `rfi` that caused the RSE references was tagged by the instruction address range or the opcode matcher. For data address range checking, the RSE reference is tagged only if its hits the programmed DBR range

## STORES\_RETIRED

- **Title:** Retired Stores, **Category:** Memory
- **Definition:** STORES\_RETIRED counts the number of retired stores. The count includes integer, FP, RSE, and uncacheable stores. Predicated-off operations are not counted
- **Event Code:** 0x6D, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

## UC\_LOADS\_RETIRED

- **Title:** Retired Uncacheable Loads, **Category:** Memory
- **Definition:** UC\_LOADS\_RETIRED counts the number of retired uncacheable or write coalescing loads. The count includes integer, FP, RSE, and VHPT loads and failed check loads (`ld.c`). Check loads that hit in the ALAT are NOT counted. Predicated-off operations are not counted
- **Event Code:** 0x6E, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

## UC\_STORES\_RETIRED

- **Title:** Retired Uncacheable Stores, **Category:** Memory
- **Definition:** UC\_STORES\_RETIRED counts the number of retired uncacheable or write coalescing stores. The count includes integer, FP, RSE, and uncacheable stores. Predicated-off operations are not counted
- **Event Code:** 0x6F, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 2
- **Qualification:** Instruction Address Range: yes, Opcode Matching: yes, Data Address Range: yes

**UNSTALLED\_BACKEND\_CYCLE**

- **Title:** Unstalled Back-end Cycles **Category:** Stall
- **Definition:** UNSTALLED\_BACKEND\_CYCLE counts the number of cycles that the back-end is processing instructions without delay and the decoupling buffer between the front-end and back-end is empty, so that any effect on the front-end will be propagated to the back-end of the pipeline. This monitor thus reflects the number of cycles where there are no back-end stalls or flushes, and the decoupling buffer is empty, regardless of whether the L1I and ITLB are being hit or missed.
- **Event Code:** 0x05, **Umask:** Ignored, **PMC/PMD:** 4, 5, 6, 7 **Max. Increment/Cycle:** 1

Instruction Address Range: no, Opcode matching: no, Data Address Range: no

# Model Specific Behavior for IA-32 Instruction Execution

The Itanium processor is capable of executing IA-32 instructions in the IA-32 system environment (legacy IA-32 operating systems) provided the required platform and firmware support exists in the system. The Itanium processor is also capable of executing IA-32 instructions in the Itanium system environment (Itanium-based operating system). Itanium-based operating system support for the capability of running IA-32 applications is defined by the respective operating system vendor. For more details on IA-32 instruction execution on Itanium-based OS, please refer to [Volume 1, Chapter 6](#) and [Volume 2, Chapter 10](#) of the *Intel® Itanium™ Architecture Software Developer's Manual*.

Note that while Itanium processor supports execution of IA-32 applications, best performance and capabilities will be realized by using 64-bit optimized OSes and applications.

In general, the behavior of IA-32 instructions on the Itanium processor is similar to that of the Pentium III processor except where noted. The following sections describe some of the key differences in behavior between IA-32 instruction execution on an Itanium processor and on the Pentium III processor. These differences do not prevent IA-32 legacy operating systems or IA-32 applications from operating correctly.

## 8.1 Processor Reset and Initialization

When RESET# is asserted, all processors based on the Itanium architecture boot at a different reset location than IA-32 processors and start executing Itanium-based 64-bit code instead of IA-32 16-bit Real Mode code. Unlike IA-32 processors, processors based on the Itanium architecture execute PAL firmware to test and initialize the processor and then continue execution in the Itanium instruction set to boot the system. SAL firmware code can switch to the IA-32 instruction set as needed to execute IA-32 BIOS code. For more details on processor reset, please refer to [Chapter 11](#) and [Chapter 13](#) of [Volume 2](#) of the *Intel® Itanium™ Architecture Software Developer's Manual*.

## 8.2 New JMPE Instruction

A new IA-32 instruction JMPE has been defined for processors based on the Itanium architecture. This instruction comes in two forms with an opcode for each. These opcodes will cause an Invalid Opcode fault on all IA-32 processors. For more details, refer to [Chapter 1](#) of [Volume 3](#) of the *Intel® Itanium™ Architecture Software Developer's Manual*.

## 8.3 System Management Mode (SMM)

SMM is superseded by the Itanium-based Platform Management definition. This mechanism is designed to provide platform level interrupt support for both IA-32 and Itanium-based operating systems. Please refer to [Chapter 11](#) of [Volume 2](#) of the *Intel® Itanium™ Architecture Software Developer's Manual* for more details on PMI.

The IA-32 SMM and I/O Port Restart feature is not supported on the Itanium processor. Dynamically, powering off/on I/O devices on an I/O Port reference via system logic is not possible for IA-32 Operating Systems or Itanium-based Operating Systems using the IA-32 SMM I/O Restart mechanism. I/O Restart has not been extended on processors based on the Itanium architecture to intercept I/O Port references from the Itanium instruction set via normal loads and stores on processors based on the Itanium architecture.

Execution of the IA-32 RSM (Resume from SMM) instruction results an Invalid Opcode fault on all processors based on the Itanium architecture.

## 8.4 CPUID Instruction Return Values for Caches and TLBs of the Intel® Itanium™ Processor

The following table provides information on how to decode return values of the CPUID instruction for the Itanium processor's internal caches and TLBs.

**Table 8-1. Encoding of Cache and TLB Return Values for the Intel® Itanium™ Processor**

Return Value	Cache or TLB Description
0x10	L0D: 16K 4-way set associative 32 bytes line
0x15	L0I: 16K 4-way set associative 32 bytes line
0x1A	L1: 96K on die 6-way set associative 64 byte line
0x88	L2: 2M 4-way set associative 64 bytes line
0x89	L2: 4M 4-way set associative 64 bytes line
0x8A	L2: 8M 4-way set associative 64 bytes line
0x90	ITLB: 4K to 256M pages, fully associative, 64 entries
0x96	DTLB0: 4K to 256M pages, fully associative, 32 entries
0x9B	DTLB1: 4K to 256M pages, fully associative, 96 entries

When the input value in register EAX is 2, the Itanium processor returns information about the processor's internal caches and TLBs in the EAX, EBX, ECX, and EDX registers. The following table describes the values returned.

**Table 8-2. EAX, EBX, ECX, and EDX Return Values for the Intel® Itanium™ Processor**

Register	Return Value (from MSB to LSB)
EAX	0x00, 0x15, 0x10, 0x01
EBX	0x00, 0x00, <L2>, 0x1A (<L2> is either 0x88, 0x89)
ECX	0x00, 0x9B, 0x96, 0x90
EDX	0x80, 0x00, 0x00, 0x00

## 8.5 Machine Check Abort (MCA)

The Itanium processor supports Pentium processor level machine checks in the IA-32 System Environment.

## 8.6 Model Specific Registers

The complete set of Model Specific Registers (MSRs) found on the Pentium III processor is not supported on the Itanium processor. For example, Model Specific Debug registers, Model Specific Test registers, Machine Check registers, and Model Specific Configuration registers are not supported.

Model Specific registers that are common to the Itanium processor and Pentium III processor use the Pentium III processor's bit definition and register assignment. The ITC, APIC\_Base, MTRR and MAP registers are supported on the Itanium processor.

## 8.7 Cache Modes

Pentium processor and Pentium III processor SRAM Cache Mode is not supported on the Itanium processor.

SRAM is typically used on IA-32 processors to provide scratch RAM areas while running IA-32 boot and machine check code before memory is available. Both of these functions are now provided by Itanium-based firmware while running IA-32 and Itanium-based operating systems.

## 8.8 10-byte Floating-point Operand Reads and Writes

Many IA-32 FP instructions read and write 10 bytes to memory. Consider the case of 16-bit segment, where the read or write starts at offset 0xFFFF8. Pentium III processor reads or writes 8 bytes then re-evaluates the linear address before reading or writing the final 2 bytes. Eight bytes are accessed at 0xffff8, and 2 bytes are accessed at 0x0000.

The Itanium processor evaluates the address once, then accesses all 10 bytes. Therefore, bytes 0xffff8 to 0x10001 will be accessed.

On a 10-byte operand read or write access, potential page faults and GP faults will return slightly different faulting addresses (linear addresses may wrap differently).

## 8.9 Floating-point Data Segment State

The Itanium processor reports a different value of the floating-point data segment state (FDS) after the execution of “FNOP” instruction (or any FP instruction that does not perform a memory reference). The contents of the data register are undefined if the prior non-control instruction did not have a memory operand. The Pentium III processor behaves as follows:

1. A FP non-transparent instruction which references memory will put the selector of the data segment used in the memory reference into FDS.
2. A FP non-transparent instruction which doesn't reference memory will put the selector of SS into FDS and 0 into FEA.

If a segment override prefix is present on an instruction of the type specified in case 2, the overriding segment selector will be put into FDS instead of the selector of SS.

The Itanium processor behavior covers only case #1 described above. Note that this difference does not affect the running of IA-32 applications.

## 8.10 Writes to Reserved Bits during FXSAVE

During FXSAVE, the Itanium processor does not write any reserved bits, while the Pentium III processor may write reserved bits. The Itanium processor does one 10 byte access to save each FP register, whereas the Pentium III processor will do two 8 byte accesses causing writes to upper reserved bits.

## 8.11 Setting the Access/Dirty (A/D) Bit on Accesses that Cross a Page Boundary

In the IA-32 system environment, the Itanium processor sets a page's A/D bit even if a memory reference crosses a page boundary and the other page has a fault. This behavior is different from Pentium III processors which do not modify the A/D bit under the above conditions.

The above difference does not come into play in the Itanium system environment.

## 8.12 Enhanced Floating-point Instruction Accuracy

On the Itanium processor, FP transcendental instructions will return more accurate (hence slightly different) answers than Pentium III processor. This behavior falls into 3 categories:

- **F2XM1, FYL2X, FYL2XP1, FPATAN Instructions**  
More accurate algorithms will result in answers which may differ from Pentium III processor



by 1 unit in the last place (ulp). Also, for FYL2X and FYL2XP1, when x or x+1 respectively is a power of two, the Precision exception is not signaled (since  $\log(2^k)$  where, k is integral, is exact).

- **FPTAN, FSIN, FCOS, FSINCOS Instructions**

New algorithms on Itanium processor include a more accurate argument reduction scheme. Although more accurate, the algorithms implemented on Itanium processor can produce answers which are different from those returned on Pentium III processor.

- **FPREM, FPREM1 Instructions**

No change.

## 8.13 RCPSS, RCPPS, RSQRTSS, RSQRTPS Instruction Differences

These four instructions are single and parallel approximations of divide and square root operations. The Itanium processor will calculate these functions to a higher accuracy than previous implementations, resulting in different answers. The Pentium III processor implementation of one of these functions can have a maximum relative error of  $1.5 * 2^{-12}$ . The Itanium processor, however, will calculate RCPPS/RCPSS functions with a maximum relative error of  $2^{-17.75288} \approx 1.1868 * 2^{-18}$  and the RSQRTPS/RSQRTSS functions with a maximum relative error of  $2^{-17.06412} \approx 1.9130 * 2^{-18}$ .

## 8.14 Read/Write Access Ordering

In general, the order of reads/writes within any complex IA-32 instruction is model specific even among IA-32 processors. Different Intel processors have different access ordering behavior; for example, internal operation ordering varies between the 80486, Pentium, Pentium III and Itanium processors.

## 8.15 Multiple IOAPIC Redirection Table Entries

If multiple IOAPIC Redirection Table Entries (RTE) share the same vector, and at least one RTE is programmed as logical delivery mode in which the selected local APIC destinations overlap with the other RTEs with the same vector, some of the selected local APICs might not receive the interrupt when the pins that correspond to these RTEs are asserted.

## 8.16 Self Modifying Code (SMC)

The Itanium processor provides the same SMC support as the Pentium processor. Also, a branch instruction is required between the store that modifies instruction(s) and the modified code.

## 8.17 Raising an Alignment Check (AC) Fault

The Pentium III processor checks and raises AC fault before a page fault. The Itanium processor checks and raise a page fault before an AC fault.

## 8.18 Maximum Number of Processors Supported in MP System Running Legacy IA-32 OS (IA-32 System Environment)

Similar to the case of IA-32 processors in an MP system, the maximum number of processors based on the Itanium architecture supported in a MP system running legacy IA-32 OS (IA-32 system environment) is 16. However, in MP systems with IA-32 processors, the number of IA-32 processors can be extended beyond 16 with additional platform enhancements while the limit for the number of processors based on the Itanium architecture running IA-32 OS in a MP system is limited to 16.

Entries in this index are described by the volume number and page or range of pages where the entries can be found. The volume number appears to the left of the colon. The page or range of pages appears to the right of the colon. A range of pages is separated by a hyphen.

## Numerics

32-bit virtual addressing 2:59, 2:60  
 pointer “swizzling” model 2:60  
 sign-extension model 2:60  
 zero-extension model 2:59, 2:60

## A

abort 2:79, 2:406  
 interruption priorities 2:92  
 machine check abort 2:44, 2:489-2:491  
 PAL-based interruptions 2:79, 2:80, 2:81, 2:85,  
 2:96, 2:406  
 PSR.mc bit is 0 2:82  
 reset abort 2:481  
 abort handling 2:491  
 access rights, segment descriptor 3:600  
 acquire semantics 1:66, 2:70, 2:238, 2:375  
 add instruction 1:45, 1:47, 1:73, 1:145, 1:177, 3:265,  
 3:352, 3:356, 3:373-3:376, 3:435, 3:448-3:450,  
 3:560, 3:620  
 address space model 2:425, 2:429  
 address translation 2:39, 2:51, 2:406, 2:425, 2:429,  
 2:485  
 addressable units 1:30  
 advanced load address table (See ALAT) 1:57, 1:62,  
 1:144, 2:416  
 ALAT 1:57-1:62, 1:144, 1:186, 2:73-2:75, 2:129, 2:130,  
 2:415, 2:416, 2:420, 2:434, 2:484, 3:306,  
 3:339, 3:346, 3:353  
 data speculation 1:57, 1:58, 1:62, 1:144, 2:445  
 related instructions 1:58  
 alloc instruction 1:14, 1:36, 1:39, 1:42-1:44, 1:48, 1:68,  
 1:136, 2:57, 2:74, 2:119, 2:121, 2:122, 2:127-  
 2:129, 3:309, 3:311, 3:339, 3:340, 3:345-3:347,  
 3:351, 3:352, 3:357  
 application programming model 1:41  
 application register  
 compare and exchange value register (CCV – AR  
 32) 1:26  
 epilog count register (EC – AR 66) 1:27  
 floating-point status register (FPSR – AR 40) 1:26  
 IA-32 time stamp counter (TSC) 1:26, 1:110, 2:27,  
 2:469  
 interval time counter (ITC – AR 44) 1:26  
 kernel registers (KR 0-7 – AR 0-7) 1:25  
 loop count register (LC – AR 65) 1:27  
 previous function state (PFS – AR 64) 1:27  
 register stack configuration register (RSC – AR 16)  
 1:25  
 RSE backing store pointer (BSP – AR 17) 1:25  
 RSE NaT collection register (RNAT – AR 19) 1:26

user NaT collection register (UNAT – AR 36) 1:26  
 application register state 1:19  
 application register model 1:21  
 ignored fields 1:20  
 ignored register 1:19, 1:20, 1:29  
 read-only register 1:20, 1:25, 2:106  
 reserved fields 1:20, 3:831  
 reserved register 1:19, 1:20  
 reserved value 1:20  
 arithmetic instructions 1:45, 3:324, 3:863  
 atomic operations 2:237  
 atomocity 2:57, 2:77, 2:238

## B

backing store pointer (BSP) 2:117, 2:121, 2:124, 2:127  
 backing store 1:24-1:26, 1:44, 2:86, 2:87, 2:117,  
 2:118, 2:485  
 backing store pointer application registers 2:124  
 backing store switches 2:130  
 BSPSTORE 2:127  
 backing store pointer application registers 2:125  
 banked general registers 2:16, 2:34, 2:35, 2:81, 2:216  
 barrier synchronization 2:394, 2:395  
 be bit 1:31  
 PSR.be 2:86, 3:343, 3:349  
 RSC.be 2:123  
 biased exponent 3:823-3:826, 3:829  
 bit field and shift instructions 1:46, 1:47  
 boot flow 2:10, 2:481  
 firmware boot flow 2:481  
 boot sequence 2:9, 2:481  
 boot flow 2:10, 2:481  
 bootstrap processor (BSP) 2:481  
 br.call instruction 1:67, 1:69, 2:121, 2:122, 2:128, 2:129,  
 2:419, 2:421, 3:318, 3:319, 3:339, 3:340,  
 3:343, 3:346, 3:347, 3:351, 3:352, 3:356, 3:357  
 br.cexit instruction 1:67, 1:69, 1:177, 3:315, 3:318,  
 3:353, 3:357  
 br.ctop instruction 1:67, 1:69, 1:177, 1:178, 1:188,  
 3:315, 3:318, 3:353, 3:357  
 br.ia instruction 1:10, 1:103, 1:105, 2:458  
 br.ret instruction 1:67, 1:69, 2:47, 2:57, 2:86, 2:92,  
 2:119, 2:121, 2:122, 2:126-2:130, 2:419, 2:421,  
 3:317, 3:319, 3:339, 3:340, 3:343, 3:345-3:347,  
 3:349, 3:351, 3:352, 3:356, 3:357  
 br.wexit instruction 1:67, 1:69, 3:315, 3:318, 3:353,  
 3:356  
 br.wtop instruction 1:67, 1:69, 1:181, 1:182, 3:315,  
 3:318, 3:353, 3:356

branch instructions 1:67, 1:70, 1:136, 1:137, 3:315, 3:316, 3:332  
    branch predict instructions 1:15, 1:70, 1:71, 3:320  
    branch prediction hints 1:70, 1:168  
    modulo-scheduled loop support 1:68  
branching 1:14, 1:22, 1:68, 1:136, 3:458, 3:528  
brl instruction 2:449  
brl.call instruction 1:67, 1:69, 2:121, 2:122, 2:128, 2:129, 2:449, 3:339, 3:340, 3:343, 3:346, 3:347, 3:351, 3:352, 3:356, 3:357  
bsw instruction 1:48, 2:18, 2:86, 3:316, 3:321, 3:343, 3:349, 3:356, 3:357  
bundles 1:11, 1:31, 1:32, 1:34, 1:132, 1:133, 3:257  
byte ordering 1:30, 1:31

## C

cache synchronization 2:69  
cache write policy attribute 2:66  
cacheability and coherency attribute 2:65  
cacheable 2:63, 2:64, 2:65, 2:66  
    cacheable pages 2:66  
    uncacheable pages 2:66  
causality 2:390  
    obeying causality 2:390  
character strings 1:76  
chk.a instruction 1:34, 1:58-1:60, 1:144-1:146, 1:148, 2:69, 2:88, 2:445, 3:305, 3:310, 3:351  
chk.a.clr instruction 1:60-1:62, 3:305, 3:310, 3:339, 3:346, 3:351  
chk.a.nc instruction 1:61, 1:62, 3:305, 3:310, 3:351  
chk.s instruction 1:34, 1:54, 1:56, 1:58, 1:134, 1:147, 1:148, 2:88, 2:445, 3:282, 3:283, 3:305, 3:311, 3:352, 3:356  
clr instruction 1:71, 3:288-3:295, 3:297-3:301, 3:305, 3:310, 3:318, 3:351-3:353  
clrrb instruction 1:36, 1:42, 1:48, 1:60, 1:68, 1:69, 2:422, 3:316, 3:321, 3:340, 3:347, 3:357  
clrrb.pr instruction 1:68, 1:69, 3:316, 3:321  
cmp instruction 1:37, 1:48, 1:50, 1:55, 1:135, 2:385, 2:394, 3:267-3:270, 3:352, 3:356, 3:418, 3:419  
cmp4 instruction 1:37, 1:48, 1:50, 1:55, 3:26-3:270, 3:352, 3:356  
cmpxchg instruction 1:26, 1:51, 1:53, 1:62, 1:66, 2:69, 2:70, 2:73, 2:179, 2:376, 2:393, 2:394, 3:351, 3:423-3:425, 3:620  
coalescing attribute 2:66  
    coalesced pages 2:67  
coherency 1:125, 2:65, 2:238, 3:576, 3:735  
compare instructions 1:14, 1:22, 1:47, 1:48, 1:182, 3:267, 3:325, 3:429  
    compare types 1:49, 1:50  
    normal compare 1:49  
    parallel compare 1:73, 1:164, 1:165, 3:327  
    unconditional compare 1:49  
computational models 1:200  
constant register 2:415  
context switching 1:43, 2:144, 2:422, 2:424, 2:477

    address space switching 2:424  
    non-local control transfer 2:422  
    performance monitor 2:144, 2:145  
    RSE backing store 1:26, 2:130  
    thread switch within the same address space 2:424  
control flow optimization 1:155  
control flow optimizations 1:163  
    multiple values for one variable or register 1:165  
    multiway branches 1:165  
    parallel compares 1:163-1:165  
control registers (CR) 2:16  
    banked general registers 2:34, 2:35, 2:81  
    control register instructions 2:25  
    default control register (DCR – CR0) 2:25, 2:26  
    external interrupt control registers 2:34, 2:104, 2:105, 2:216, 2:467  
    global control registers 2:25  
    interruption control registers 2:29  
    interruption faulting address (IFA – CR20) 2:31  
    interruption function state (IFS – CR23) 2:33  
    interruption hash address (IHA – CR25) 2:34  
    interruption immediate (IIM – CR24) 2:34  
    interruption instruction bundle pointer (IIP – CR19) 2:30  
    interruption instruction previous address (IIPA – CR22) 2:32, 2:33  
    interruption processor status register (IPSR – CR16) 2:29  
    interruption status register (ISR – CR17) 2:29  
    interruption status register fields 2:29  
    interruption TLB insertion register (ITIR – CR21) 2:32  
    interruption vector address (IVA – CR2) 2:28  
    interval time counter (ITC – AR44) 2:27  
    interval timer match register (ITM – CR1) 2:27  
    ITIR fields 2:32  
    page table address (PTA – CR8) 2:28  
control registers, moving values to and from 3:636  
control speculative 1:12, 2:415, 2:445  
corrected machine check (CMC) 2:109, 2:283, 2:489, 2:490  
corrected machine check interrupt (CMCI) 2:489  
cover instruction 1:36, 1:42, 1:44, 1:48, 1:69, 2:87, 2:121, 2:122, 2:125-2:127, 2:129, 3:316, 3:321, 3:339, 3:340, 3:344, 3:346, 3:347, 3:351, 3:357  
CUID registers 1:29  
cross-modifying code 2:400  
current frame marker (CFM) 1:19, 1:23, 1:36, 1:42, 2:13, 2:121  
    size of frame (sof) 1:42  
    size of locals (sol) 1:42  
cycles per instructions (CPI) 2:477

## D

data access bit 2:57, 2:59, 2:90, 2:91  
data breakpoint register matching 2:246  
    DBR.addr 2:246  
    DBR.mask 2:247  
    trap code B bits 2:247  
data breakpoint registers (DBR) 2:133, 2:134

data debug 2:57, 2:90, 2:91  
 data dependencies 1:57, 1:140-1:142, 2:13, 2:378, 2:383  
 data dependency 1:14, 1:141, 1:142, 2:13, 2:382-2:384  
 data key miss fault 2:128, 2:149  
 data key permission 2:57, 2:59, 2:90, 2:91  
 data NaT page consumption 2:57, 2:59, 2:72, 2:89-2:91  
 data nested TLB faults 2:59, 2:413  
 data page not present 2:57, 2:59, 2:89, 2:91  
 data prefetch  
   load instructions 1:52  
   semaphore instructions 1:53, 3:303  
   store instructions 1:52, 3:287, 3:301, 3:303, 3:304, 3:924  
 data serialization 2:14, 2:15, 3:337  
 data speculative 1:13, 1:57, 1:146, 2:415  
 data TLB miss faults 2:57, 2:59  
 debug 1:59, 2:177  
   break instruction fault 2:133, 2:149, 2:165  
   data debug fault 2:94, 2:128, 2:134, 2:149, 2:177  
   debug breakpoint registers (DBR/IBR) 2:16  
   debug instructions 2:135  
   debug model 2:246  
   debugging 2:133, 2:462, 3:597  
   debugging facilities 2:133  
   instruction breakpoints 2:462  
   instruction debug fault 2:133, 2:149, 2:177, 2:246  
   lower privilege transfer trap 2:133, 2:182, 2:246  
   single step trap 2:84, 2:94, 2:97, 2:133, 2:150, 2:151, 2:184, 2:246, 2:462  
   taken branch trap 2:84, 2:94, 2:97, 2:133, 2:150, 2:151, 2:183, 2:246, 2:462, 3:598  
 Dekker's algorithm 2:396  
 dependencies 1:35-1:38, 1:143, 3:335, 3:336  
   dependency violation 1:36, 1:38, 2:448  
   instruction execution 3:427, 3:551, 3:576, 3:681, 3:735  
   instruction group 1:35-1:39, 1:67  
   register dependencies 1:35-1:39  
   WAR dependency 1:38, 3:350  
 division operations  
   double precision – divide 1:199  
   double precision – square root 1:199  
 DMA 1:17, 2:403  
   edge sensitive interrupt messages 2:114

## E

edge- and level-sensitive interrupts 2:114  
 EFI 2:253, 2:451, 2:481-2:483, 2:485, 2:488  
   boot services 2:483, 2:488  
   EFI boot manager 2:482  
   EFI procedure calls 2:488  
   EFI system partition 2:482, 2:483  
   runtime services 2:488  
 EFLAGS register  
   condition codes 3:415, 3:458, 3:463  
   flags affected by instructions 3:366  
   loading 3:599

popping 3:665  
 popping on return from interrupt 3:579  
 pushing 3:673  
 pushing on interrupts 3:565  
 saving 3:695  
   status flags 3:418, 3:588, 3:704, 3:729  
 exception deferral 1:55, 2:45, 2:90, 2:91  
   combined hardware/software deferral 2:443, 2:444  
   exception deferral of control speculative loads 2:443  
   hardware-only deferral 2:443, 2:444  
   software-only deferral 2:443, 2:444  
 exception indicator 2:72  
 exception qualification 2:89  
 execution unit type 1:24, 1:32, 1:33, 3:257  
 extended instructions 1:34, 3:257  
 extensible firmware interface (See EFI) 2:451, 2:481  
 external interrupt (INT) 2:92, 2:101  
   control register usage examples 2:467  
   external (I/O) devices 2:98  
   external interrupt (INT)  
     external interrupt architecture 2:463  
   external interrupt delivery 2:81, 2:102-2:104, 2:106, 2:465  
   external interrupt masking 2:103, 2:464  
   external interrupt sampling 2:104  
   external interrupt states 2:101  
   inactive 2:465  
   in-service/none pending 2:466  
   in-service/one pending 2:466  
   internal processor interrupts 2:99  
   interrupt acknowledge (INTA) cycle 2:114  
   interrupt enabling 2:102  
   interrupt masking 2:103  
   interrupt priorities 2:102, 2:463  
   interrupt registers 2:16  
   interrupt sources 2:98, 2:104, 2:114, 2:464  
   interrupt vectors 2:101-2:104, 2:464, 3:565  
   locally connected devices 2:98  
   pending 2:82, 2:99, 2:102, 2:103, 2:465, 2:466  
 external interrupt control registers 2:34, 2:104, 2:105, 2:216, 2:464, 2:467  
   Local ID (LID – CR64) 2:105  
 external task priority (XTP) 2:110, 2:114  
   XTP cycle 2:114  
   XTP register 2:465

## F

fault 1:95, 2:26, 2:81, 2:82, 2:85, 2:86, 2:92-2:96, 2:128, 2:405, 2:449  
 fault suppression 2:88  
 fclass instruction 1:37, 1:48, 1:50, 1:55, 1:92, 1:202, 3:325, 3:326, 3:352, 3:356  
 fcmp instruction 1:37, 1:48, 1:50, 1:55, 1:91, 3:325, 3:326, 3:339, 3:351, 3:356  
 fence 1:66, 2:69, 2:70  
   operations 1:66, 2:69, 2:70  
   semantics 1:66, 2:69, 2:70, 2:238, 2:239

- fetchadd instruction 1:51, 1:53, 1:62, 1:66, 2:69, 2:70, 2:73, 2:179, 2:376, 2:377, 2:394, 2:395, 3:351
  - firmware address space 2:258
  - firmware entrypoint 2:283
  - firmware entrypoints 2:257
  - firmware interface table (FIT) 2:261
  - firmware model 2:482
  - firmware procedure 2:283
  - floating-point applications 1:193
    - execution bandwidth 1:194
    - execution latency 1:193
    - memory bandwidth 1:195
    - memory latency 1:194
    - performance limiters 1:193
  - floating-point architecture 1:11, 1:15, 1:77
  - floating-point instructions 1:26, 1:37, 1:83, 3:322, 3:323, 3:366, 3:489, 3:491, 3:510, 3:526, 3:543, 3:748
    - arithmetic instructions 1:90, 3:324, 3:863
    - integer multiply and add instructions 1:93, 1:94
    - memory access instructions 1:83
    - non-arithmetic instructions 1:92
    - register to/from general register transfer instructions 1:89
  - floating-point programming model 1:77
    - data types and formats 1:77
    - floating-point register encodings 1:78
    - floating-point register format 1:77, 1:78
    - floating-point status register 1:24, 1:26, 1:80, 1:81, 1:93
    - real types 1:77
  - floating-point status register (FPSR) 1:26, 1:37, 1:80, 1:93, 2:451
  - floating-point system software
    - floating-point exception handling 2:451, 2:453
  - flushrs instruction 1:36, 1:44, 1:48, 2:57, 2:119, 2:121, 2:122, 2:125-2:130, 3:306, 3:311, 3:339, 3:340, 3:345, 3:346, 3:357
  - FP precision 1:197
  - FP subfield handling 1:203
  - FPU
    - checking for pending FPU exceptions 3:734
    - constants 3:487
    - existence of 3:428
    - floating-point format 3:822, 3:823
    - initialization 3:481
    - saving 3:510, 3:526
    - storing 3:524
  - FPU data pointer 3:491, 3:508, 3:510, 3:526
  - FPU flag, CPUID instruction 3:428
  - FPU instruction pointer 3:491, 3:508, 3:510, 3:526
  - FPU last opcode 3:491, 3:508, 3:510, 3:526
  - FPU status word
    - condition code flags 3:460, 3:476, 3:536, 3:538, 3:541
    - FPU flags affected by instructions 3:366
    - loading 3:491
    - restoring 3:508
    - saving 3:510, 3:526, 3:528
    - TOP field 3:480
  - FPU tag word 3:491, 3:508, 3:510, 3:526
  - frcpa instruction 1:37, 1:48-1:50, 1:55, 1:91, 1:198-1:200, 2:454, 3:323, 3:326, 3:352, 3:356
  - frsqta instruction 1:37, 1:48-1:50, 1:55, 1:91, 1:198, 1:199, 2:454, 3:323, 3:327, 3:352, 3:356
- ## G
- gate interception 2:215
  - general register (GR)
    - NaT bit 1:21, 1:134, 1:146, 1:147
  - global TLB purge operations 2:69
- ## H
- hardware debugger 2:136
- ## I
- i bit
    - PSR.i 2:82, 2:102-2:104, 2:106, 2:219, 2:407, 2:410, 2:464, 2:465, 3:344, 3:349, 3:410, 3:721
  - I/O port space 2:240-2:243, 2:473, 2:474, 3:558, 3:562, 3:654, 3:657
  - I/O port space model 2:240, 2:241
    - physical I/O port addressing 2:243
    - virtual I/O port addressing 2:241
  - IA-32 application execution model 1:103
    - IA-32 instruction set execution 1:22, 1:23, 1:41, 1:59, 1:104, 2:239
    - IA-32 operating mode transitions 1:106
    - instruction set execution in the Itanium architecture 1:104
    - instruction set modes 1:103
    - instruction set transitions 1:105, 2:215, 2:240
  - IA-32 application register state model 1:107
    - IA-32 application EFLAG register 1:116
    - IA-32 floating-point registers 1:117, 1:118
    - IA-32 general purpose registers 1:107, 1:108, 1:110
    - IA-32 instruction pointer 1:111
    - IA-32 MMX technology registers 1:122
    - IA-32 segment registers 1:111

- IA-32 streaming SIMD extension registers 1:109, 1:122
- IA-32 application support 2:250
  - procedure calls between Itanium and IA-32 instruction sets 2:459
  - transitioning between Itanium and IA-32 instruction sets 2:457
- IA-32 architecture 1:5, 1:17, 2:5, 3:5, 3:359
- IA-32 architecture handlers 2:460
  - IA-32 vectors that need Itanium-based OS support 2:461
  - shared Itanium/IA-32 exception vectors 2:460
  - unique IA-32 exception vectors 2:460
  - unique Itanium exception vectors 2:460
- IA-32 compatible bus transactions 2:251
- IA-32 current privilege level 2:218
- IA-32 fault and trap handling 2:215
- IA-32 faults 3:359
- IA-32 floating-point exceptions 2:456
- IA-32 GPFault 3:359
- IA-32 I/O instructions 2:244
- IA-32 instruction behavior 2:215, 2:227
- IA-32 instruction format 3:360
- IA-32 instruction summary 2:228
- IA-32 interruption 2:94, 2:95, 2:248
- IA-32 interruption priorities and classes 2:95
- IA-32 interruption vector 2:189, 2:248
- IA-32 memory ordering 2:238, 2:392
- IA-32 MMX technology instructions 1:122, 3:747
- IA-32 numeric exception model 2:250
- IA-32 physical memory references 2:235
- IA-32 privileged system resources 2:215
- IA-32 processes during a context switch 2:226
  - entering IA-32 processes 2:226
  - exiting IA-32 processes 2:227
- IA-32 segmentation 1:124, 2:233
- IA-32 streaming SIMD extension instructions 1:123, 3:811
- IA-32 system and control register behavior 2:215
- IA-32 system EFLAG register 2:219
- IA-32 system environment 1:5, 1:9, 1:10, 1:17, 2:5, 2:9, 3:5
- IA-32 system register mapping 2:216
- IA-32 system registers 2:222
  - IA-32 control registers 2:222
  - IA-32 debug registers 2:225
  - IA-32 machine check registers 2:226
  - IA-32 memory type range registers (MTRRs) 2:225
  - IA-32 model specific and test registers 2:225
  - IA-32 performance monitor registers 2:226
- IA-32 system segment registers 2:217
- IA-32 TLB forward progress requirements 2:234
- IA-32 trap code 2:189
- IA-32 usage of Itanium registers 1:126
  - ALAT 1:126
  - NaT/NaTVal response for IA-32 instructions 1:126
  - register stack engine 1:126, 3:597
- IA-32 virtual memory references 2:234
  - protection keys 2:234
  - region identifiers 2:234
  - TLB access bit 2:234
  - TLB dirty bit 2:234
- IA-32 virtual memory support 2:215
- ic bit
  - PSR.ic 2:82, 2:85, 2:86, 2:88-2:91, 2:102, 2:104, 2:126, 2:148, 2:407, 2:430, 2:431, 3:344, 3:349
- IEEE considerations 1:94
  - additions beyond the IEEE standard 1:100
  - arithmetic operations 1:100, 3:826
  - floating-point interruptions 1:94
  - inexact 1:97, 1:99, 2:453, 2:456
  - integer invalid operations 1:100
  - mandated operations deferred to software 1:100
  - NaNs 1:78, 1:100
  - overflow 1:97, 1:98, 2:452, 2:455
  - tininess 1:99
  - underflow 1:97, 1:99, 2:452, 2:456
- IEEE floating-point exception filter 2:451, 2:454
  - denormal/unnormal operand exception (fault) 2:455
  - divide by zero exception (fault) 2:455
  - inexact exception (trap) 2:456
  - invalid operation exception (fault) 2:455
  - overflow exception (trap) 2:455
  - underflow exception (trap) 2:456
- IEEE-754 2:451, 2:454, 2:456, 3:821
  - ANSI/IEEE-754 standard compliant 1:193
- if-conversion 1:157
- illegal dependency fault 1:38, 2:149, 2:448
- illegal operation fault 1:19, 1:20, 1:39, 2:149, 3:257
- implicit serialization 2:13
- indefinite
  - description of 3:827
  - real 3:829
- infinity, floating-point format 3:826
- in-flight resources 2:15
- INIT flows 2:491
- initialization event (INIT) 2:489
  - initialization interrupts 2:80, 2:98, 2:406
  - PALE\_INIT 2:80, 2:92
- inserting/purging of translations 2:425
- instruction breakpoint register matching 2:247
  - IBR.addr 2:247
  - IBR.mask 2:247
- instruction breakpoint registers (IBR) 2:133, 2:134
- instruction classes 3:336, 3:350, 3:351
- instruction dependencies 1:140
  - control dependencies 1:65, 1:140, 2:384
  - data dependencies 1:57, 1:141-1:143, 2:13
- instruction encoding 1:32
  - bundles 1:11, 1:31, 1:32, 1:34, 1:132, 1:133, 3:257
  - instruction slots 1:32, 1:33, 3:257
  - template 1:32-1:34, 1:133, 3:257, 3:258
- instruction field names 3:259, 3:262
- instruction format 1:132, 3:260
- instruction interception 2:215
- instruction pointer (IP) 1:19, 1:22, 2:84, 2:408, 2:479

- instruction serialization 2:14, 2:15, 2:410, 3:337, 3:431, 3:579
  - instruction set
    - string instructions 3:420, 3:562, 3:622, 3:640, 3:656, 3:724
  - instruction set architecture (ISA) 1:5, 2:5, 3:5
  - instruction set features 1:11
  - instruction set transition model overview 1:10
  - instruction set transitions 2:33, 2:215, 2:240
  - instruction slots 1:32, 1:33, 3:257
    - instruction slot mapping 1:33, 3:258
  - instruction stream 1:167, 3:496, 3:510, 3:526, 3:591, 3:649
    - instruction stream alignment 1:167
    - instruction stream fetching 1:167
  - instruction type 1:32, 3:257, 3:697
    - ALU (A) 3:258
    - branch (B) 3:258
    - floating-point (F) 3:258
    - integer (I) 1:133, 3:258
    - memory (M) 1:133, 3:258
  - instruction/data TLB miss 2:57-2:59
  - integer computation instructions 1:44
    - 32-bit addresses and integers 1:46
    - arithmetic instructions 1:45, 3:324, 3:863
    - bit field and shift instructions 1:46, 1:47
    - large constants 1:47
    - logical instructions 1:45
  - integer/floating-point conversion 1:203
  - inter-processor interrupt (IPI) 2:97, 2:99, 2:110, 2:471
  - inter-processor interrupt message 2:111, 2:491
    - data fields 2:111, 2:113
  - interrupt 2:70, 2:80, 2:81, 2:92, 2:96-2:106, 2:405, 2:406, 2:463, 2:464, 2:491
  - interrupt acknowledge (INTA) 2:110
  - interruption 2:79, 2:80-2:87, 2:92, 2:94-2:96, 2:103, 2:104, 2:129, 2:405-2:407
    - execution environment 2:407
    - heavyweight interruptions 2:411, 2:413
    - interruption handler 2:86, 2:87, 2:405-2:407, 2:410
    - interruption handling 2:79, 2:82, 2:85, 2:86, 2:410
    - interruption register state 2:408
    - lightweight interruptions 2:410
    - nested interruptions 2:413
    - resource serialization 2:409, 2:410
  - interruption model 2:82, 2:247
  - interruption priorities 2:92, 2:95
  - interruption registers 2:216, 2:406, 2:408
  - interruption vector address (IVA) 2:406
  - interruption vector table (IVT) 2:79, 2:96, 2:406
  - interruption vectors 2:85, 2:96, 2:147, 2:151, 2:406
    - interruption vector definition 2:148
  - interruptions 2:79-2:82, 2:85-2:87, 2:92, 2:95, 2:96, 2:127, 2:128, 2:405, 2:406
    - aborts 2:79, 2:89, 2:92, 2:406
    - faults 2:79, 2:80, 2:85, 2:89, 2:92-2:96, 2:405
    - interruption handling during instruction execution 2:82
    - interruption programming model 2:81
    - interruptions 2:79, 2:80, 2:82, 2:85, 2:89, 2:92, 2:95, 2:97-2:106, 2:219, 2:405, 2:406
    - IVA-based interruption 2:85, 2:96, 2:406
    - PAL-based interruption 2:85, 2:405
    - traps 1:97, 2:79, 2:80, 2:84-2:86, 2:92, 2:94-2:96, 2:405
  - interval timer 1:110, 2:16, 2:27, 2:28, 2:99, 2:108, 2:469, 2:470
  - invala instruction 1:60, 1:62, 2:129, 2:420, 2:484, 3:306, 3:307, 3:311, 3:351, 3:352
  - invala.e instruction 1:60, 1:62, 2:415-2:417, 2:420, 3:306, 3:307, 3:311, 3:339, 3:351, 3:352
  - IPI ordering 2:113
  - ISR setting 2:147
  - Itanium architecture 1:1, 1:5, 1:9
  - Itanium data mem faults 3:360
  - Itanium instruction 1:107, 3:257, 3:335, 3:430, 3:597, 3:598
    - expressing parallelism 1:132
    - format 3:464, 3:539
    - Itanium instruction set 1:17, 3:430, 3:597, 3:598
    - syntax 1:132, 3:336
  - Itanium instruction mem faults 3:360
  - Itanium system environment 1:5, 1:9, 1:10, 1:17, 2:5, 2:9, 2:10, 2:11
  - Itanium-based firmware 1:5, 1:17, 2:5, 3:5
  - itc instruction 1:24, 1:26, 1:27, 2:27, 2:43, 2:44, 2:47, 2:50, 2:58, 2:429, 2:431, 2:469, 2:470, 3:312, 3:313, 3:339-3:342, 3:344-3:346, 3:348, 3:352-3:354, 3:356
  - itr instruction 2:40, 2:43, 2:44, 2:47, 2:50, 2:429, 2:430, 2:485, 2:488, 3:312, 3:313, 3:340-3:342, 3:344, 3:345, 3:348, 3:352, 3:356
  - IVA-based interruptions 2:79-2:81, 2:85, 2:405, 2:406
- ## J
- jmpe instruction 1:10, 1:103, 1:105
- ## L
- Lampport's algorithm 2:397, 2:398
  - ld.a instruction 1:51, 1:57, 1:61, 1:62, 1:144, 1:146, 1:153, 2:68, 2:69, 2:74, 2:376, 2:443-2:445
  - ld.acq instruction 1:51, 1:60, 1:66, 2:69, 2:70, 2:376, 2:381, 2:383, 2:385, 2:387, 2:389, 2:390
  - ld.c instruction 1:57-1:60, 1:144-1:146, 1:153, 2:69, 2:444, 2:445
  - ld.c.clr instruction 1:51, 1:60-1:62, 2:73, 2:74
  - ld.c.clr.acq instruction 1:51, 1:60-1:62, 1:66, 2:69, 2:70, 2:73, 2:74
  - ld.c.nc instruction 1:51, 1:61, 1:62, 2:73, 2:74
  - ld.s instruction 1:51, 1:54, 1:56, 1:148, 2:68, 2:69, 2:376, 2:443-2:445
  - ld.sa instruction 1:51, 1:61, 1:62, 1:148, 2:68, 2:69, 2:74, 2:376, 2:416, 2:443-2:445



ld8.fill instruction 1:26, 1:36, 1:37, 1:52, 1:55, 1:56,  
1:147, 2:415, 2:416, 3:288, 3:289, 3:293-3:295,  
3:338, 3:340, 3:352

ldf.a instruction 1:51, 1:57, 1:61, 1:62

ldf.c instruction 1:57

ldf.c.clr instruction 1:51, 1:61, 1:62, 2:73

ldf.c.nc instruction 1:51, 1:61, 1:62, 2:73

ldf.fill instruction 1:51, 1:52, 1:56, 1:83, 1:147, 2:69,  
2:415, 2:416, 3:290, 3:291, 3:297-3:299, 3:352

ldf.s instruction 1:51, 1:54, 1:56, 2:69

ldf.sa instruction 1:51, 1:61, 1:62, 2:69

ldfp.a instruction 1:51, 1:57, 1:59, 1:61, 1:62

ldfp.c instruction 1:57

ldfp.c.clr instruction 1:51, 1:61, 1:62

ldfp.c.nc instruction 1:51, 1:61, 1:62

ldfp.s instruction 1:51, 1:54, 1:56, 2:69

ldfp.sa instruction 1:51, 1:61, 1:62, 2:69

level sensitive external interrupts 2:114

load instruction 1:159, 2:376, 3:873, 3:876, 3:879, 3:882,  
3:884

loadrs instruction 1:25, 1:36, 1:44, 1:48, 2:57, 2:92,  
2:119, 2:121-2:123, 2:125-2:130, 2:418, 2:419,  
2:485, 3:306, 3:311, 3:339, 3:340, 3:345,  
3:346, 3:351, 3:357

loadrs field 1:44, 2:122, 2:125  
RSC.loadrs 1:44, 2:125-2:127, 2:485

logical instructions 1:45

long branch handler 2:447

loop support 1:68, 1:171, 1:174  
capacity limitations 1:185  
conflicts in the ALAT 1:186  
counted loop 1:69, 1:171, 1:176, 1:177  
counted loop branches 1:176  
epilog 1:68, 1:174, 1:179  
epilog count register (EC) 1:27  
explicit prolog and epilog 1:190  
implementing reductions 1:189  
induction variable 1:172  
initiation interval (II) 1:173  
kernel 1:68, 1:174, 1:175, 1:179  
kernel iteration 1:174  
kernel loop 1:174  
loop count application register (LC) 1:69, 1:171  
loop unrolling 1:137, 1:172, 1:187  
loop unrolling prior to software pipelining 1:187  
loops with predicated instructions 1:182  
multiple-exit loops 1:183  
prolog 1:68, 1:174, 1:179  
redundant load elimination in loops 1:192  
register rotation 1:15, 1:174, 1:175  
software pipelining and advanced loads 1:185  
software pipelining considerations 1:185  
software-pipelined loop branches 1:176, 1:177  
source iteration 1:174  
source loop 1:174  
while loop 1:70, 1:178, 1:180, 3:686, 3:687

**M**

machine check 2:43, 2:44, 2:79, 2:92, 2:224, 2:226,  
2:283, 2:406, 2:489-2:491, 3:429, 3:679, 3:737

machine check abort (See MCA)

machine check abort flows  
machine check abort handling in OS 2:491  
machine check handling in PAL 2:490  
machine check handling in SAL 2:490

machine check aborts 2:481

machine checks 2:268

major opcode 1:34, 3:257-3:259

master boot record 2:483

mc bit  
PSR.mc 2:82, 2:85, 2:86, 2:102, 2:408, 3:345,  
3:349

MCA 2:481  
PALE\_CHECK 2:79, 2:92

memory acceptance fence 2:473

memory access control 1:204  
allocation control 1:63, 1:205  
data prefetch 1:205  
load-pair instructions 1:204

memory access instructions 1:51, 1:62, 2:375

memory access ordering 1:65, 2:70  
memory ordering instructions 1:66  
memory ordering rules 1:66

memory addressing model 1:30, 1:123

memory alignment 2:236

memory attribute 2:45, 2:63, 2:64, 2:73-2:75, 3:429  
effects of memory attributes on advanced/check  
loads 2:73  
effects of memory attributes on memory reference  
instructions 2:72  
memory attribute transition 2:74  
physical addressing memory attribute 2:64  
virtual addressing memory attribute 2:63, 2:74

memory dependency 1:35, 2:69  
read-after-write 1:35, 1:38, 1:65, 2:69  
write-after-read 1:35, 1:38, 1:65, 2:69  
write-after-write 1:35, 1:38, 1:65, 2:69

memory endianess 1:124

memory fence 1:66, 2:392

memory fences 2:113, 2:378

memory hierarchy 1:63  
hierarchy control and hints 1:62  
memory consistency 1:65, 3:924-3:926

memory mapped I/O model 2:241, 2:474

memory model 1:123, 2:233

memory ordering 1:66, 2:69, 2:130, 2:375, 2:376, 2:378,  
2:379, 2:383-2:385, 2:424, 3:928  
acquire semantics 1:66, 2:70, 2:238, 2:375  
memory ordering executions 2:379  
memory ordering interactions 1:125  
memory ordering model 2:238, 2:378, 2:392  
memory ordering semantics 1:66, 2:378, 2:424  
release semantics 1:66, 2:69, 2:238, 2:239, 2:375

memory ordering fence 1:66

memory reference 1:139, 1:140, 2:38

memory synchronization 2:392  
 mf instruction 1:66, 2:113, 2:378  
 mf.a instruction 2:69, 2:113, 2:473, 2:474, 2:475, 3:306, 3:311  
 Min/Max/AMin/AMax 1:202  
 MMX technology 1:16, 1:104, 1:107, 1:109, 1:122, 3:359, 3:429, 3:747-3:749  
 mov instruction 1:29, 1:37, 1:38, 1:47, 1:50, 1:68, 1:74, 1:75, 2:13, 2:18, 2:19, 2:47, 2:50, 2:124, 2:127, 2:135, 2:140, 2:141, 2:477, 3:282-3:286, 3:307-3:309, 3:311-3:314, 3:320, 3:353-3:355, 3:388, 3:618, 3:632-3:639, 3:659, 3:660, 3:717  
 multimedia instructions 1:11, 1:16, 1:41, 1:72  
   data arrangement 1:74  
   parallel arithmetic 1:72, 1:73  
   parallel shifts 1:73  
 multimedia support 1:16  
 multiple address space (MAS) 1:16, 2:37, 2:425, 2:426  
 multiple status fields 1:200  
 multiply-add instruction 1:198  
 multiprocessor (MP)  
   multiprocessor instruction cache coherency 2:238  
   multiprocessor TLB coherency 2:235

## N

NaN  
   description of 3:824, 3:826  
   encoding of 3:825, 3:829  
   operating on 3:827  
   SNaNs vs. QNaNs 3:826  
   testing for 3:536  
 NaNs 1:78, 1:100, 1:202, 3:453, 3:497, 3:499, 3:502, 3:531, 3:534, 3:547, 3:549, 3:822, 3:824-3:827, 3:829  
 NaT (not a thing) 1:131  
 NaT page consumption fault 2:72  
 NaTPage attribute 2:72  
 NaTVal (not a thing value) 1:22  
 non-access instructions 2:87  
 non-cacheable memory 2:69  
 non-programmer-visible state 2:378  
 non-speculative 1:54, 2:67, 2:68, 2:74, 2:445  
 non-speculative memory references 1:139, 2:62  
   data prefetch hint 1:140  
   loads from memory 1:139  
   stores to memory 1:139  
 non-temporal hint 1:205  
 NOP instruction 3:649  
 no-recovery model 2:88, 2:89  
 normalized finite number 3:823, 3:825  
 normalized numbers 1:79, 3:822, 3:823, 3:825  
 not a thing attribute (NaTPage) 2:72

## O

OLR 2:283  
 operand screening support 1:202

operating environments 1:9, 1:10  
 optimization of memory references 1:148  
   data interference 1:149, 1:150  
   loop optimization 1:152  
   minimizing check code 1:152  
   optimizing code size 1:150  
   using post-increment loads and stores 1:151  
 orderable instruction 2:376, 2:380  
 ordered cacheable operations 2:389  
 ordering semantics 1:66, 2:69, 2:70, 2:380  
   acquire 1:66, 2:69, 2:70, 2:238, 2:239, 2:375, 2:380, 2:381  
   fence 1:66, 2:69, 2:70, 2:239, 2:375, 2:380  
   release 1:66, 2:69, 2:70, 2:238, 2:239, 2:375, 2:380, 2:381  
   unordered 1:66, 2:69, 2:70, 2:239, 2:375, 2:380  
 OS boot flow sample code 2:495  
 OS kernel 2:483, 2:485  
 OS loader 2:482, 2:483  
 overflow 1:14, 1:97, 1:98, 2:452, 2:455

## P

PAL 1:5, 1:17, 2:5, 2:253, 2:256, 2:283, 2:481, 2:483, 2:485-2:490, 3:5  
   entrypoints 2:256  
   procedures 2:256  
 PAL power on/reset 2:263  
   PALE\_RESET 2:79  
 PAL procedure calling conventions 2:288  
 PAL procedure calls 2:486  
 PAL procedures 2:285, 2:481, 2:486-2:488  
   stacked PAL call 2:487  
   stacked registers 1:136, 2:486, 2:487  
   static PAL call 2:486  
 PAL self-test control word 2:267  
 PAL\_MC\_RESUME 2:276  
 PAL\_PREFETCH\_VISIBILITY 2:354  
 PAL\_PROC\_GET\_FEATURES 2:355  
 PAL-based interrupt states 2:100  
 PAL-based interruptions 2:79-2:81, 2:85, 2:96, 2:405, 2:406  
 PALE\_CHECK 2:268  
 PALE\_INIT 2:277  
 PALE\_RESET 2:263  
 performance counters 1:28, 2:137, 2:138, 2:225, 2:477, 3:681  
 performance monitor events 2:142  
 performance monitors 1:28, 2:137, 2:139, 2:217, 2:477, 2:478  
   performance monitor code sequences 2:143  
   performance monitor configuration (PMC) 2:137, 2:139  
   performance monitor data (PMD) 2:137, 2:477  
   performance monitor data registers (PMD) 1:19, 1:28  
   performance monitor interrupt service routine 2:143  
   performance monitor overflow registers 2:141

- performance monitor registers 2:137, 2:139, 2:477
  - performance monitoring mechanisms 2:477
  - physical addressing 2:61, 2:64, 2:76, 2:488, 3:429, 3:637
  - pk bit 2:427
    - PSR.pk 2:83, 2:86, 2:427, 2:428, 3:345, 3:349
  - platform management interrupt (PMI) 2:92, 2:279, 2:405, 2:481, 2:489, 2:492
  - PMI flows 2:492
  - population count 1:76, 3:279
  - power management 2:80, 2:281, 2:492
    - NORMAL 1:164, 3:541, 3:905, 3:907
  - predicate register (PR)
    - predicate register transfers 1:50
  - predication 1:11, 1:13, 1:47, 1:48, 1:135, 1:138, 1:155-1:160
    - cache pollution reduction 1:160
    - downward code motion 1:159, 1:160
    - guidelines for removing branches 1:162
    - instruction prefetch hints 1:168
    - instruction scheduling 1:140, 1:142, 1:156
    - off-path predication 1:158
    - optimizing program performance using predication 1:157
    - performance costs of branches 1:155
    - predication considerations 1:160
    - predication in the Itanium architecture 1:156
    - prediction resources 1:71, 1:155, 1:156
    - upward code motion 1:159
  - preservation of floating-point state in the OS 2:419
  - preserved 2:283
  - preserved registers 2:415, 2:420
  - preserving ALAT coherency 2:420
  - privilege levels 1:22, 2:13, 3:597, 3:598, 3:688
    - current privilege level (CPL) 2:13, 3:732
    - privilege level transfer 1:76
    - processor status register (PSR) 2:13, 2:16, 2:18
    - processor status register fields 2:19
    - processor status register instructions 2:18
  - privileged operation fault 2:149
  - probe instruction 2:47, 2:50, 2:62, 2:63, 2:87, 3:312, 3:313, 3:356
  - procedure 1:41-1:43
  - procedure calls 1:41, 1:136, 2:415, 2:459, 2:485-2:488
    - br.call 1:67, 1:69, 3:318, 3:319
    - br.ret 1:67, 1:69, 2:47, 2:57, 2:86, 2:92, 3:317, 3:319
    - branch instructions 1:70, 1:136, 1:137, 3:315, 3:316, 3:332
    - branches and hints 1:136
    - loops and software pipelining 1:137
    - register stack engine 1:41, 1:136, 2:86, 2:117, 3:597
    - rotating registers 1:23, 1:137
    - stacked register 1:42, 2:486, 2:487
  - processor abstraction layer (See PAL)
  - processor caches 2:75, 2:378
  - processor identifiers (CPUID) 1:19
    - processor identification registers 1:29
  - processor interrupt block 2:110-2:112, 2:471
  - processor min-state save area 2:274
  - processor ordered 2:238
  - processor state 2:289
    - system state 2:13, 2:15, 2:16
  - processor state parameter 2:271
  - processor status register (PSR) 2:13, 2:16, 2:18, 2:139, 2:408
  - programmed I/O 2:401, 2:402
  - protected mode 1:10, 1:104, 1:106, 1:113-1:115, 2:224, 2:458
  - protection key registers (PKR) 2:16, 2:48
  - protection keys 1:16, 2:16, 2:48, 2:49, 2:425, 2:427-2:429, 2:435
  - pseudo-code functions 3:249
  - ptc.e instruction 2:41, 2:43, 2:50, 2:57, 2:426, 2:432-2:434, 3:307, 3:312, 3:341, 3:342, 3:348, 3:352, 3:356
  - ptc.g instruction 2:41, 2:43, 2:44, 2:47, 2:50, 2:57, 2:63, 2:69, 2:433, 2:434, 3:312, 3:314, 3:341, 3:342, 3:345, 3:348, 3:352, 3:356
  - ptc.ga instruction 1:60, 2:41, 2:43, 2:44, 2:47, 2:50, 2:57, 2:63, 2:69, 2:426, 2:433, 2:434, 3:312, 3:314, 3:341, 3:342, 3:345, 3:348, 3:352, 3:356
  - ptr instruction 2:40, 2:43, 2:47, 2:50, 2:57, 2:429, 2:431, 2:485, 3:312, 3:314, 3:341, 3:342, 3:345, 3:348, 3:352, 3:356, 3:571
- ## Q
- QNaN
    - description of 3:826
    - operating on 3:827
  - qualified exception deferral 2:91
- ## R
- RAR (read-after-read) dependency 3:335
  - RAW (read-after-write) dependency 3:335
  - reader of a resource 3:335
  - real mode 1:10, 1:104, 1:106, 1:113-1:115, 2:458, 3:491, 3:508, 3:510, 3:526, 3:579, 3:665
  - recovery model 2:88, 2:89
  - region identifier (RID) 2:38, 2:48, 2:425
  - region register (RR) 2:48, 2:425
  - register dependency 1:35, 1:37
    - read-after-write (RAW) 1:35
    - write-after-read (WAR) 1:35
    - write-after-write (WAW) 1:35
  - register file transfers 1:74
  - register preservation 2:415
    - preservation at different points in the OS 2:418
    - preservation of stacked registers in the OS 2:418
    - preserving floating-point registers 2:416
    - preserving general registers 2:416
  - register rotation 1:15, 1:23, 1:174, 1:175
    - initializing rotating predicates 1:50, 1:176

- register stack 1:14, 1:23, 1:24, 1:41-1:44, 2:87, 2:117, 2:119-2:121, 2:485
    - clean partition 2:120, 2:127
    - current frame 1:23, 1:42, 2:86, 2:87, 2:117, 2:120, 2:485
    - dirty partition 2:120, 2:127
    - invalid partition 2:120, 2:127
    - register stack instructions 1:43
    - register stack operation 1:41
  - register stack configuration 1:24, 1:25, 1:44, 2:119, 2:122, 2:123, 2:485
    - RSC 1:24, 1:25, 1:44, 2:119, 2:122, 2:123, 2:127, 2:485, 3:353, 3:355
  - register stack engine (See RSE)
  - release semantics 1:66, 2:69, 2:238, 2:239, 2:375
  - release stores 2:376, 2:378, 2:389, 2:390
  - reserved 1:19, 1:20, 2:97, 2:284
  - rfi instruction 1:11, 1:34, 1:36, 1:38, 1:48, 1:69, 1:103, 2:18, 2:19, 2:57, 2:63, 2:79, 2:81, 2:84, 2:86-2:88, 2:92, 2:119, 2:121, 2:122, 2:127-2:130, 2:182-2:184, 2:406, 2:410, 2:420, 2:421, 2:485, 3:316, 3:321, 3:339-3:341, 3:343-3:347, 3:349-3:351, 3:356, 3:357
  - RSE 1:25, 1:26, 1:41, 1:136, 2:86, 2:117-2:119, 2:121-2:123, 2:125-2:131, 2:485, 3:345, 3:350, 3:351, 3:357
    - RSE byte order 2:123
    - RSE control instructions 2:125, 2:126
    - RSE initialization 2:132
    - RSE internal state 2:119
    - RSE interruptions 2:127
    - RSE mode 1:25, 2:122
    - RSE operation instructions and state modification 2:122
    - RSE privilege level 1:25, 2:123
  - rsm instruction 2:18, 2:19, 2:33, 2:104, 2:140, 2:232, 2:410, 2:464, 2:477, 3:311, 3:314, 3:341, 3:356, 3:357, 3:694
  - rum instruction 1:75, 2:13, 2:18, 2:140, 2:477, 3:311, 3:314, 3:338, 3:345, 3:357
- ## S
- SAL 1:5, 1:17, 2:5, 2:253, 2:284, 2:458, 2:481-2:483, 2:485-2:491, 3:5, 3:696, 3:697, 3:698, 3:699, 3:709
    - SAL procedure calls 2:487
  - SALE\_ENTRY 2:265
  - scratch 2:284
  - scratch registers 2:81, 2:415, 2:420
  - self test state parameter 2:266
  - self-modifying code 2:399
  - semaphore 3:303
    - behavior of uncacheable and misaligned semaphores 2:377
    - semaphore instructions 1:35, 1:53, 2:376, 3:303
    - semaphore operations 1:53, 2:237, 2:378, 2:388
  - sequential consistency (SC)
    - SC system 2:392
  - sequential semantics 2:70
    - inter-processor interrupt messages 2:70, 2:111-2:113
    - sequential pages 2:70
  - serialization 2:13-2:15, 2:409, 2:410, 3:335-3:337, 3:431, 3:579, 3:597
  - single address space (SAS) 1:16, 2:37, 2:425, 2:427, 2:429
  - single instruction multiple data (SIMD) 3:812
  - single stepping 2:88
  - sof field
    - CFM.sof 2:84, 2:120-2:122, 2:126, 2:127, 2:129
  - software pipelining 1:11, 1:15, 1:137, 1:173, 1:185, 1:187
  - sol field
    - CFM.sol 2:122, 2:127, 2:129
  - special instruction notations 3:263
  - special use registers 2:415
  - speculation 1:11, 1:12, 1:133, 1:134, 1:139, 1:143, 1:149, 2:67, 2:68, 2:443, 2:445
    - advanced load 1:51, 1:57-1:59, 1:61, 1:144-1:146, 1:153, 2:68, 2:73, 2:74
    - advanced load check 1:58, 1:59, 1:146, 3:305
    - advanced load example 1:145
    - always-defer model 2:88
    - check load 1:51, 1:57, 1:59-1:61, 1:145, 1:146, 2:73, 2:74
    - combining data and control speculation 1:148
    - control speculation 1:12, 1:13, 1:53-1:56, 1:61, 1:134, 1:143, 1:146, 1:147, 2:445
    - control speculation example 1:147
    - control speculative load 1:12, 1:146-1:148
    - data speculation 1:12, 1:13, 1:57, 1:58, 1:61, 1:62, 1:134, 1:143, 1:144, 2:445
    - recovery code 1:13, 1:58, 1:145, 1:146-1:148, 2:444, 2:445
    - recovery code example 1:145
    - speculation attributes 2:67
    - speculation check 1:54, 1:58, 1:148, 3:283, 3:305
    - speculation considerations 1:149
    - speculation model in the itanium architecture 1:143, 1:144
    - speculation recovery code 2:445
    - speculation related exception handlers 2:445
    - speculative 1:12, 1:13, 1:54, 1:57, 2:67, 2:68, 2:74, 2:445
    - speculative load exceptions 2:89
    - speculatively accessible 2:68
    - speculatively inaccessible 2:68
    - unaligned handler 2:445
  - speculative advanced load 1:148
  - spill/fill 1:56, 1:83, 1:89, 2:86, 2:117, 2:118, 2:121
  - spin lock 2:393, 2:394
  - square root operations 1:198
  - ssm instruction 2:18, 2:19, 2:33, 2:103, 2:104, 2:140, 2:410, 2:477, 3:311, 3:314, 3:341, 3:356, 3:357, 3:694

- st instruction 1:13, 1:51, 1:62, 2:69, 2:380-2:385, 2:387, 2:390, 2:391, 3:357, 3:445, 3:447, 3:448, 3:451, 3:453, 3:456, 3:458, 3:460, 3:461, 3:463, 3:464, 3:466, 3:469, 3:472, 3:475, 3:476, 3:478, 3:482, 3:483, 3:485, 3:487, 3:493, 3:497-3:500, 3:502, 3:503, 3:505, 3:507, 3:508, 3:510, 3:513-3:515, 3:517, 3:519, 3:521, 3:530, 3:533, 3:536, 3:538, 3:539, 3:541, 3:543, 3:545, 3:547-3:549
  - st.rel instruction 1:51, 1:62, 1:66, 2:69, 2:113, 2:376, 2:381, 2:383-2:385, 2:388-2:391
  - st.spill instruction 1:51, 1:62, 2:69
  - st1 instruction 1:60, 3:288, 3:289, 3:296, 3:357, 3:856, 3:859, 3:860
  - st8.spill instruction 1:26, 1:36-1:38, 1:53, 1:55, 1:56, 1:147, 2:415, 2:416, 3:288, 3:289, 3:296, 3:338, 3:340, 3:347, 3:357
  - stack frame 1:14, 1:23, 1:25, 1:36, 1:41-1:44, 2:117, 2:119, 3:442-3:444
  - stacked calling convention 2:284
  - stacked registers 1:21, 1:42, 1:136, 2:117, 2:118, 2:120, 2:415, 2:418, 2:485
    - deallocated 2:129
    - stacked general registers 1:21, 2:117, 2:416
  - state mappings 3:359
  - static calling convention 2:284
  - static general registers 1:21, 2:117, 2:416
  - stf instruction 1:51, 1:62, 2:69, 3:290, 3:291, 3:299, 3:300, 3:357
  - stf.spill instruction 1:51, 1:53, 1:56, 1:62, 1:83, 1:147, 2:69, 2:415, 2:416, 3:290, 3:291, 3:299, 3:300, 3:357
  - store buffers 2:378, 2:385, 2:387
  - store instruction 2:376, 3:873, 3:876, 3:879, 3:882, 3:884, 3:925, 3:926, 3:928
  - streaming SIMD extension technology 1:104, 3:359, 3:430
  - subpaging 2:440, 2:441
  - sum instruction 1:75, 2:13, 2:18, 2:140, 2:477, 3:311, 3:314, 3:338, 3:345, 3:356, 3:357, 3:435, 3:448, 3:739, 3:910, 3:919, 3:920
  - supervisor accesses 2:236
  - system abstraction layer (See SAL)
  - system architecture features 1:16, 2:11
    - support for multiple address space operating systems 1:16
    - support for single address space operating systems 1:16
    - system performance and scalability 1:17
    - system security and supportability 1:17
  - system calls 2:420, 2:421, 2:422
  - system descriptors 2:217
  - system flag interception 2:215
  - system memory model 2:233
  - system register model 2:17, 2:215
    - IA-32 state 1:108, 2:215, 2:216
    - shared 1:108, 2:216, 2:217, 3:620
    - undefined 1:108, 2:216
    - unmodified 1:108, 1:109, 2:216, 3:582
  - system register resources 2:13, 2:15, 2:16
- ## T
- tak instruction 2:49, 2:50, 2:63, 2:87, 2:428, 3:312, 3:315, 3:341, 3:345, 3:352, 3:356
  - tbit instruction 1:37, 1:48, 1:50, 1:55, 1:135, 3:281, 3:352, 3:356
  - template 1:32, 1:33, 1:34, 1:133, 3:257, 3:258
  - temporal hint 1:205, 3:927
  - thash instruction 2:50, 2:54-2:56, 2:438, 2:439, 3:312, 3:315, 3:340, 3:341, 3:345, 3:352, 3:356
  - TLB 1:59, 2:17, 2:28, 2:32, 2:37-2:59, 2:96, 2:234, 2:235, 2:485
    - page not present vector 2:97, 2:152, 2:440
    - TLB miss 2:42, 2:51, 2:52, 2:55-2:59, 2:436
    - TLB miss handlers 2:59, 2:436, 2:439
    - TLB purges 2:40, 2:42, 2:44
    - translation insertion format 2:45
    - VHPT translation vector 2:96, 2:152, 2:438
  - TLB entry, invalidating (flushing) 3:578
  - tnat instruction 1:37, 1:48, 1:50, 1:55, 1:56, 3:281, 3:282, 3:352, 3:356
  - tpa instruction 2:50, 2:62, 2:87, 3:312, 3:315, 3:341, 3:345, 3:352, 3:356
  - translation caches (TCs) 2:431
    - TC insertion 2:431
    - TC purge 2:429, 2:432
  - translation lookaside buffer (See TLB)
  - translation registers (TRs) 2:429
    - TR insertion 2:430
    - TR purge 2:429, 2:430, 2:431
  - trap instruction 1:97, 1:98, 2:84, 2:85, 2:94, 2:96, 2:405, 3:411, 3:565-3:569, 3:573-3:575, 3:581, 3:585, 3:597, 3:598, 3:601, 3:627, 3:633, 3:634, 3:659, 3:661, 3:667, 3:674, 3:721, 3:722
  - ttag instruction 2:50, 2:53-2:56, 2:437, 3:312, 3:315, 3:345, 3:352, 3:356
- ## U
- UC memory attribute 2:242
  - unaligned reference handler 2:445-2:447
  - uncacheable 2:63-2:65, 3:927
    - uncacheable pages 2:66
  - unchanged 2:19, 2:159, 2:284, 3:466, 3:481, 3:505, 3:513, 3:515, 3:517, 3:666, 3:688, 3:865, 3:867, 3:869, 3:871, 3:875, 3:876, 3:878, 3:879
  - undefined behavior 1:39
  - underflow 1:14, 1:97, 1:99, 2:452, 2:456
  - unimplemented addresses 2:62, 2:63
    - unimplemented physical address bits 2:61, 2:62
    - unimplemented virtual address bits 2:62
  - unnormalized numbers 1:79
  - unordered semantics 2:375
  - unsupported data reference handler 2:447, 2:448
  - user mask (UM) 1:19, 1:28

## V

vector numbers 2:80, 2:101, 2:463, 3:367, 3:573  
VHPT 2:28, 2:34, 2:37-2:39, 2:41, 2:47, 2:48, 2:50-2:59,  
2:96, 2:434-2:436, 2:485  
    TLB and VHPT search faults 2:59  
    TLB/VHPT search 2:58  
    translation searching 2:57  
    VHPT configuration 2:51  
    VHPT searching 2:52  
    VHPT short format 2:52  
    VHPT short-format index 2:54, 2:55  
    VHPT updates 2:436  
    VHPT walker 2:39, 2:41, 2:48, 2:51-2:59, 2:434-  
    2:437  
virtual addressing 2:37, 2:38, 2:63, 2:74, 2:485, 2:488  
virtual aliasing 2:60  
virtual hash page table (See VHPT)  
virtual region number (VRN) 2:38, 2:62, 2:425  
virtualized interrupt flag 2:219  
visible 1:66, 2:69, 2:70, 2:376, 2:382, 3:600, 3:626, 3:928  
VM86 1:10, 1:104, 1:106, 1:113, 1:114, 2:221, 2:224,  
2:458, 3:571, 3:572

VME extensions 2:219, 2:224

## W

WAR (write-after-read) dependency 3:335  
WAW (write-after-write) dependency 3:335  
write BSPSTORE 2:131  
write-back and invalidate caches 3:735  
writer of a resource 3:335

## X

xchg instruction 1:51, 1:53, 1:62, 1:66, 2:69, 2:70, 2:73,  
2:179, 2:376, 2:388, 3:357, 3:387, 3:620, 3:741,  
3:742

## Z

zero, floating-point format 3:824