



# Revision Guide for AMD Athlon™ 64 and AMD Opteron™ Processors

Advanced Micro Devices 

Publication # <b>25759</b>	Revision: <b>3.43</b>
Issue Date: <b>April 2005</b>	

© 2003–2005 Advanced Micro Devices, Inc. All rights reserved.

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

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

#### **Trademarks**

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

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

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

## Revision History

---

<b>Date</b>	<b>Revision</b>	<b>Description</b>
April 2005	3.43	Added erratum #121–122.
November 2004	3.39	Added AMD Sempron™ processor and AMD Athlon™ 64 FX processor (939 package) information; Added erratum #117; Changed first sentence in Suggested Workaround for erratum #81 and #97.
June 2004	3.27	Added erratum #112.
June 2004	3.25	Added erratum #109 and #111; Updated CPUID values in Table 1; Updated Documentation Support section.
April 2004	3.23	Added erratum #108; Updated CPUID value in Table 1.
April 2004	3.21	Added erratum #105–107.
February 2004	3.19	Added errata #103–104; Changed Suggested Workaround section in erratum #1; Updated CPUID values in Table 1; Updated die information in Table 2; Added package names to Table 3; Added Mixed Silicon Revision Support information to Revision Determination section on page 6.
December 2003	3.15	Added erratum #101.
October 2003	3.11	Added errata #99–100. Added reference document to Suggested Workaround section of erratum #98.
September 2003	3.09	Added AMD Athlon™ 64, Mobile AMD Athlon 64, and AMD Athlon 64 FX information; Revised erratum #91. Corrected erratum #77.
August 2003	3.07	Added erratum #95, #97, and #98; Modified Description in erratum #94; Added C0 silicon information.
June 2003	3.05	Added erratum #94.
May 2003	3.03	Added erratum #92.
April 2003	3.01	Initial public release.



---

# Revision Guide for AMD Athlon™ 64 and AMD Opteron™ Processors

---

The purpose of the *Revision Guide for AMD Athlon™ 64 and AMD Opteron™ Processors* is to communicate updated product information to designers of computer systems and software developers. This revision guide includes information on the following products:

- AMD Athlon™ 64 processor
- Mobile AMD Athlon™ 64 processor – includes desktop replacement (DTR)
- AMD Athlon™ 64 FX processor
- AMD Opteron™ processor
- AMD Sempron™ processor

This guide consists of three major sections:


- **Revision Determination:** This section, starting on page 6, describes the mechanism by which the current revision of the part is identified.
- **Product Errata:** This section, starting on page 7, provides a detailed description of product errata, including potential effects on system operation and suggested workarounds. An erratum is defined as a deviation from the product's specification, and as such may cause the behavior of the processor to deviate from the published specifications.
- **Documentation Support:** This section, starting on page 65, provides a listing of available technical support resources.

## Revision Guide Policy

Occasionally, AMD identifies product errata that cause the processor to deviate from published specifications. Descriptions of identified product errata are designed to assist system and software designers in using the processors described in this revision guide. This revision guide may be updated periodically.

# Revision Determination

Figure 1 shows the format of the value returned in EAX by CPUID Function 1.



**Figure 1. Format of CPUID Value Returned by Function 1**

Table 1 shows the identification number returned by the CPUID instruction for each revision of the processor.

**Table 1. CPUID Values for Revisions of the Processors**

Revision	CPUID Function 1 EAX Value				
	AMD Opteron™ Processor	AMD Athlon™ 64 Processor	AMD Athlon™ 64 FX Processor	Mobile AMD Athlon™ 64 FX Processor	AMD Sempron™ Processor
SH7-B3	00000F51h	N/A	N/A	N/A	N/A
SH7-C0	00000F58h	00000F48h	00000F58h	00000F48h	N/A
SH7-CG	00000F5Ah	00000F4Ah (754) 00000F7Ah (939)	00000F5Ah (940) 00000F7Ah (939)	00000F4Ah	N/A
DH7-CG	N/A	00000FC0h (754) 00000FE0h (754) 00000FF0h (939)	N/A	00000FC0h 00000FE0h	00000FC0h (754) 00000FE0h (754) 00000FF0h (939)
CH7-CG	N/A	00000F82h (754) 00000FB2h (939)	N/A	00000F82h	N/A

## Mixed Silicon Revision Support

AMD Opteron processors with different silicon revisions can be mixed in a multiprocessor system as described in the *BIOS and Kernel Developer's Guide for AMD Athlon™ 64 and AMD Opteron™ Processors*, order# 26094. Mixed revision support includes only rev. C0 and CG AMD Opteron processors. Refer to Table 1 for CPUIDs for these revisions. Errata workarounds must be applied according to revision as described in the Product Errata section starting on page 7.

## Product Errata

This section documents product errata for the processors. A unique tracking number for each erratum has been assigned within this document for user convenience in tracking the errata within specific revision levels. Table 2 cross-references the revisions of the part to each erratum. An “X” indicates that the erratum applies to the revision. The absence of an “X” indicates that the erratum does not apply to the revision. An “\*” indicates advance information that the erratum has been fixed but not yet verified. “No fix planned” indicates that no fix is planned for current or future revisions of the processor.

**Note:** *There may be missing errata numbers. Errata that have been resolved from early revisions of the processor have been deleted, and errata that have been reconsidered may have been deleted or renumbered.*

**Table 2. Cross-Reference of Product Revision to Errata**

No.	Errata Description	Revision Number				
		SH7-B3	SH7-C0	SH7-CG	DH7-CG	CH7-CG
1	Inconsistent Global Page Mappings Can Lead to Machine Check Error	No fix planned				
51	Arbitrated Interrupt With Illegal Vector Sets APIC Error Bit In All Processors	X	X	X	X	X
57	Some Data Cache Tag Eviction Errors Are Reported As Snoop Errors	No fix planned				
58	Memory Latency with Processor Power States	X				
60	Single Machine Check Error May Report Overflow	No fix planned				
61	Real Mode RDPMC with Illegal ECX May Cause Unpredictable Operation	X				
62	Task Gates With Breakpoints Enabled May Cause Unexpected Faults	X				
63	TLB Flush Filter Causes Coherency Problem in Multiprocessor Systems	X				
64	Single Step Across I/O SMI Skips One Debug Trap	X				
65	Uncorrectable NB Machine Check Error May Disrupt Power Management	X				
66	Upstream Read Response Delayed by Downstream Posted Writes	X				
68	Disable DQS Hysteresis Bit Not Readable	X				
69	Multiprocessor Coherency Problem with Hardware Prefetch Mechanism	X				
71	RSM from SMM with Paging Enabled May Deadlock	X				

**Table 2. Cross-Reference of Product Revision to Errata (Continued)**

No.	Errata Description	Revision Number				
		SH7-B3	SH7-C0	SH7-CG	DH7-CG	CH7-CG
74	Registered DIMM Exit-Self-Refresh Requirements Not Met	X				
75	APIC Timer Accuracy Across Power Management Events	No fix planned				
76	APIC Timer Undercounts In Divide-by-8 Low Power Mode	X				
77	Long Mode CALLF or JMPF May Fail To Signal GP When Callgate Descriptor is Beyond GDT/LDT Limit	No fix planned				
78	APIC Interrupt Latency With C2 Enabled	X	X			
79	Power Management Limitations Above 1.50V	X	X	X	X	X
80	Registered DIMM Initialization Requirements Not Met	X				
81	Cache Coherency Problem with Hardware Prefetching and Streaming Stores	X				
82	Certain Faults On Far Transfer Instructions In 64-Bit Mode Save Incorrect RIP	X				
83	DC Machine Check Extended Error Code Bit Not Writeable	X	X	X	X	X
85	Insufficient Delay Between MEMCLK Startup and CKE Assertion During Resume From S3		No fix planned			
86	DRAM Data Masking Feature Can Cause ECC Failures		X	X	X	X
88	SWAPGS May Fail To Read Correct GS Base	X	X			
89	Potential Deadlock With Locked Transactions	No fix planned				
90	False IC Machine Check Overflow Error Logged On Reset	X	X	X	X	X
91	Software Prefetches May Report A Page Fault	X	X			
92	Deadlock In Multi-Processor Systems May Occur When Earlier Operations Prevent An Older Store From Writing Data	X	X			
93	RSM Auto-Halt Restart Returns To Incorrect RIP		X			
94	Sequential Prefetch Feature May Cause Incorrect Processor Operation	X	X	X	X	X
95	RET Instruction May Return To Incorrect EIP	X	X			
96	Increased Memory Latency During P-State Changes	X	X	X	X	X
97	128-Bit Streaming Stores May Cause Coherency Failure		X	X	X	X
98	LDTSTOP Assertion May Be Missed		X	X	X	X
99	Background Scrubbing Must Be Disabled With Non-Contiguous Memory Map	X	X	X	X	X
100	Compatibility Mode Branches Transfer to Illegal Address	X	X			
101	DRAM Scrubber May Cause Data Corruption When Using Node-Interleaved Memory	X	X	X	X	X
103	AAM or DIV followed by AAM May Produce Incorrect Results	X	X	X	X	X



**Table 2. Cross-Reference of Product Revision to Errata (Continued)**

No.	Errata Description	Revision Number				
		SH7-B3	SH7-C0	SH7-CG	DH7-CG	CH7-CG
104	DRAM Data Masking Feature Causes ChipKill ECC Failures When Enabled With x8/x16 DRAM Devices		X	X	X	X
105	Misaligned 128-bit Store May Cause Deadlock	X	X	X	X	X
106	Potential Deadlock with Tightly Coupled Semaphores in an MP System	X	X	X	X	X
107	Possible Multiprocessor Coherency Problem with Setting Page Table A/D Bits	X	X	X	X	X
108	CPUID Instruction May Return Incorrect Model Number In Some Processors				X	
109	Certain Reverse REP MOVES May Produce Unpredictable Behavior		X	X	X	X
111	Rtt Specification Violation			X	X	X
112	Self-Modifying Code May Execute Stale Instructions	X	X	X	X	X
117	Incorrect Value May Be Returned When Reading ChipKill ECC Syndrome	X	X	X	X	X
121	Sequential Execution Across Non-Canonical Boundary Causes Processor Hang	X	X	X	X	X
122	TLB Flush Filter May Cause Coherency Problem in Multiprocessor Systems		X	X	X	X

Table 3 cross-references the errata to each processor segment. An empty cell signifies that the erratum does not apply to the processor segment. “X” signifies that the erratum applies to the processor segment. “N/A” signifies that the erratum does not apply to the processor segment due to the silicon revision.

**Table 3. Cross-Reference of Errata to Processor Segments**

Errata Number	AMD Opteron™ Processor (940 Package)	AMD Athlon™ 64 Processor (754/939 Packages)	Mobile AMD Athlon™ 64 Processor (754 Package)	AMD Athlon™ 64 FX Processor (940 Package)	AMD Athlon™ 64 FX Processor (939 Package)	AMD Sempron™ Processor (754 and 939 Packages)
1	X	X	X	X	X	X
51	X					
57	X	X	X	X	X	X
58	X	N/A	N/A	N/A	N/A	N/A
60	X	X	X	X	X	X
61	X	N/A	N/A	N/A	N/A	N/A
62	X	N/A	N/A	N/A	N/A	N/A
63	X	N/A	N/A	N/A	N/A	N/A
64	X	N/A	N/A	N/A	N/A	N/A
65	X	N/A	N/A	N/A	N/A	N/A
66	X	N/A	N/A	N/A	N/A	N/A
68	X	N/A	N/A	N/A	N/A	N/A
69	X	N/A	N/A	N/A	N/A	N/A
71	X	N/A	N/A	N/A	N/A	N/A
74	X	N/A	N/A	N/A	N/A	N/A
75	X	X	X	X	X	X
76	X	N/A	N/A	N/A	N/A	N/A
77	X	X	X	X	X	X
78	X	X	X	X	N/A	N/A
79	X	X	X	X	X	X
80	X	N/A	N/A	N/A	N/A	N/A
81	X	N/A	N/A	N/A	N/A	N/A
82	X	N/A	N/A	N/A	N/A	N/A
83	X	X	X	X	X	X
85	X	X	X	X	X	X
86	X	X	X	X	X	X

**Table 3. Cross-Reference of Errata to Processor Segments (Continued)**

Errata Number	AMD Opteron™ Processor (940 Package)	AMD Athlon™ 64 Processor (754/939 Packages)	Mobile AMD Athlon™ 64 Processor (754 Package)	AMD Athlon™ 64 FX Processor (940 Package)	AMD Athlon™ 64 FX Processor (939 Package)	AMD Sempron™ Processor (754 and 939 Packages)
88	X	X	X	N/A	N/A	N/A
89	X	X	X	X	X	X
90	X	X	X	X	X	X
91	X	X	X	X	X	X
92	X					
93	X	X	X	N/A	N/A	N/A
94	X	X	X	X	X	X
95	X	X	X	X	N/A	N/A
96		X	X			X
97	X	X	X	X	X	X
98	X	X	X	X	X	X
99	X	X	X	X	X	X
100	X	X	X	X	N/A	N/A
101	X					
103	X	X	X	X	X	X
104	X					
105	X	X	X	X	X	X
106	X					
107	X					
108		X	X			X
109	X	X	X	X	X	X
111	X	X	X	X	X	X
112	X	X	X	X	X	X
117	X					
121	X	X	X	X	X	X
122	X					

# **1 Inconsistent Global Page Mappings Can Lead to Machine Check Error**

## **Description**

If the same linear to physical mapping exists in multiple CR3 contexts, and that mapping is marked global in one context and not global in another context, then a machine check error may be reported by the TLB error detection logic (depending on the specific access pattern and TLB replacements encountered).

## **Potential Effect on System**

In the somewhat unlikely event that all required conditions are present (including the effects of the TLB replacement policy), then an unexpected machine check error may be reported. If the erratum occurs in the instruction cache TLB (L1 or L2), the apparent error is logged and corrected. If the erratum occurs in the data cache TLB (L1 or L2), the apparent error is logged and reported as an uncorrectable machine check error.

## **Suggested Workaround**

None required. This is not expected to occur in real systems.

## **Fix Planned**

No

## 51 Arbitrated Interrupt With Illegal Vector Sets APIC Error Bit In All Processors

### Description

If an arbitrated interrupt uses an illegal interrupt vector (0–15), then the corresponding error bit is erroneously set in the APIC Error Status Register (ESR) of all processors, not just the processor that accepted the arbitrated interrupt. The accepting processor has its error bit set twice—once during the arbitration phase (at the same time as all the other processors are erroneously setting their error bits), and once during the acceptance phase.

### Potential Effect on System

If an arbitrated interrupt uses an illegal vector, all processors report that error in their APIC ESR. This is not expected to impact system operation.

In the case of the accepting processor, the error may appear to be reported twice if software clears the ESR in the short interval between the arbitration and acceptance phases.

### Suggested Workaround

None required.

### Fix Planned

Yes

## **57 Some Data Cache Tag Eviction Errors Are Reported As Snoop Errors**

### **Description**

In some cases, the machine check error code on a data cache (DC) tag array parity error erroneously classifies an eviction error as a snoop error.

The common cases of cache line replacements and external probes are classified correctly (as eviction and snoop respectively). The erroneous cases occur when a tag error is detected during a DC eviction that was generated by a hardware prefetch, a cache line state change operation, or a number of other internal microarchitectural events. In such cases, the error code logged in the DC Machine Check Status register (MC0\_STATUS, MSR 0x401) erroneously indicates a snoop error.

### **Potential Effect on System**

Internally detected DC tag errors may be reported to software as having been detected by snoops. Depending upon machine check software architecture, the system response to such errors may be broader than necessary.

### **Suggested Workaround**

None required.

### **Fix Planned**

No

## **58 Memory Latency with Processor Power States**

### **Description**

If CPU Low Power mode is enabled in the C1, C2, or throttling processor power states, then externally generated sequences of memory references may experience unexpectedly large latencies through the memory controller.

### **Potential Effect on System**

Long memory latencies may lead to performance anomalies or functional failures, depending on the buffering capabilities of external devices.

### **Suggested Workaround**

Do not enable CPU Low Power mode in the C1, C2, or throttling processor power states. Specifically, disable the CPU`LowPwrEn` bits for System Management Action Field (SMAF) codes 000, 101, and 111 by clearing `Dev:3x80[0]` for C2, `Dev:3x84[24]` for C1, and `Dev:3x84[8]` for throttling.

### **Fix Planned**

Yes

## **60 Single Machine Check Error May Report Overflow**

### **Description**

A single parity error encountered in the data cache tag array may incorrectly report the detection of multiple errors, as indicated by the overflow bit of the DC Machine Check Status register (bit 62 of MSR 0x401).

### **Potential Effect on System**

System software may be informed of a machine check overflow when only a single error was actually encountered.

### **Suggested Workaround**

Do not rely on the state of the OVER bit in the DC Machine Check Status register.

### **Fix Planned**

No



## **61 Real Mode RDPMC with Illegal ECX May Cause Unpredictable Operation**

### **Description**

Illegal values of ECX (that is,  $ECX > 3$ ) for the RDPMC (Read Performance Monitor Counter) instruction correctly cause the processor to take a general protection exception.

However, if the RDPMC instruction is executed in real mode with a specific illegal value of  $ECX = 9$ , then the processor may incorrectly enter the GP fault handler as if it were in 32-bit mode.

### **Potential Effect on System**

Incorrect instruction decode leading to unpredictable system failure.

### **Suggested Workaround**

When in real mode, restrict use of the RDPMC instruction to the legal counter values (0–3). This circumstance is not expected to occur in normal operation and has only been detected in a simulation environment.

### **Fix Planned**

Yes

## **62 Task Gates With Breakpoints Enabled May Cause Unexpected Faults**

### **Description**

When a task gate is used by a CALL or JMP instruction and any debug breakpoint is enabled through the DR7.LE or GE bits, the processor may incorrectly use the new TSS base [15:0] contained in the new TSS as a selector. This will most likely lead to a GP fault with an error code of the new TSS base.

### **Potential Effect on System**

Unexpected faults leading to unpredictable system failure.

### **Suggested Workaround**

When running software that uses task gates with CALL or JMP instructions, do not enable debug breakpoints.

### **Fix Planned**

Yes

## **63 TLB Flush Filter Causes Coherency Problem in Multiprocessor Systems**

### **Description**

If the TLB flush filter is enabled in a multiprocessor configuration, coherency problems may arise between the page tables in memory and the translations stored in the on-chip TLBs. This can result in the possible use of stale translations even after software has performed a TLB flush.

### **Potential Effect on System**

Unpredictable system failure.

### **Suggested Workaround**

In MP systems, disable the TLB flush filter by setting HWCR.FFDIS (bit 6 of MSR 0xC001\_0015).

### **Fix Planned**

Yes

## **64 Single Step Across I/O SMI Skips One Debug Trap**

### **Description**

When single stepping (with EFLAGS.TF) across an IN or OUT instruction that detects an SMI, the processor correctly defers taking the debug trap and instead enters SMM. Upon RSM (without I/O restart), the processor should immediately enter the debug trap handler.

Under this scenario, the processor does not enter the debug trap handler but instead returns to the instruction following the I/O instruction.

### **Potential Effect on System**

When using the single step debug mode, following an I/O operation that detects an SMI, one instruction may appear to be skipped.

### **Suggested Workaround**

None required as this is a debug limitation only. If a workaround is desired, modify the SMM handler to detect this case and enter the debug handler directly.

### **Fix Planned**

Yes

## 65 Uncorrectable NB Machine Check Error May Disrupt Power Management

### Description

If an uncorrectable machine check error in the Northbridge (NB) error reporting bank is detected at approximately the same time as any of the following events, then the intended power management activity may be disrupted and various failures may result:

- A HyperTransport™ link frequency change
- Entry into the S3 state (suspend-to-RAM)

### Potential Effect on System

For HyperTransport link frequency changes, undefined operation results, leading to unpredictable system failure.

For entry into the S3 power state, the DRAM is not put into self-refresh state, leading to likely corruption of DRAM contents.

### Suggested Workaround

A workaround may not be required since this erratum only occurs in the presence of a fatal machine check error, and even then only when it happens to coincide in a small window of time with one of the power management events described.

If desired, the NB error reporting bank MCG\_CTL, MSR 0x17B, bit 4 (NBE), (or all of the machine check architecture) can be disabled by software around the time of these events. Refer to the *BIOS and Kernel Developer's Guide for AMD Athlon™ 64 and AMD Opteron™ Processors*, order# 26094, for information on how to program the machine check architecture.

### Fix Planned

Yes

## **66 Upstream Read Response Delayed by Downstream Posted Writes**

### **Description**

An upstream read to main memory can be delayed when the following sequence occurs:

1. The processor issues one or more posted writes downstream.
2. The processor evicts a line from its cache.
3. The chipset performs an upstream read to memory with the PassPW bit set in the HyperTransport packet.

In this case, the read should pass the downstream posted writes but due to a resource conflict in the internal request queues, the read is delayed until the processor's cache line is written and all previously enqueued posted writes have completed.

### **Potential Effect on System**

Unexpectedly large latencies may be experienced during upstream memory reads, potentially resulting in performance anomalies or functional failures, depending on the buffering capabilities of external devices.

### **Suggested Workaround**

Limit the number of processor downstream posted requests to one by programming the following register values:

- Set Dev:3x70[1:0] to 10b (SRI-to-XBAR buffer counts)
- Set Dev:3x7C[5:4] to 00b (Free list buffer counts)

Note that a warm reset is required to allow these new values to take effect.

### **Fix Planned**

Yes

## **68 Disable DQS Hysteresis Bit Not Readable**

### **Description**

The DisDqsHys bit of the DRAM Config Low register (Dev:2x90[3]) is writeable and performs its intended function, but it incorrectly always reads as zero to software.

### **Potential Effect on System**

None expected. This bit is configured by BIOS and readability is not required.

### **Suggested Workaround**

Perform writes to the DisDqsHys bit as appropriate for the system, disregarding the value read back.

### **Fix Planned**

Yes

## **69 Multiprocessor Coherency Problem with Hardware Prefetch Mechanism**

### **Description**

If the on-chip hardware prefetch mechanism generates a prefetch with write intent for a cache line that is also found to be present in the instruction cache, then the eventual prefetch response from the system is incorrectly discarded by the processor.

In the event the prefetched line was transferred in the modified state from another processor's cache, that processor's modified data is lost.

### **Potential Effect on System**

Multiprocessor memory coherency issues leading to unpredictable system failure.

### **Suggested Workaround**

In MP systems, set BU\_CFG.WbPfSmcChkDis (bit 45 of MSR 0xC001\_1023). No loss of performance results from this workaround.

### **Fix Planned**

Yes



## **71 RSM from SMM with Paging Enabled May Deadlock**

### **Description**

Under a rare set of internal timing circumstances, a speculative TLB reload may incorrectly interact with the RSM instruction such that the processor becomes deadlocked. This can only occur if the SMM handler configures and enables its own paging environment.

### **Potential Effect on System**

The system hangs and recovers only after a system reset is performed.

### **Suggested Workaround**

If paging is enabled in the SMM handler, disable it (by clearing CR0.PG) before executing the RSM instruction.

### **Fix Planned**

Yes

## **74 Registered DIMM Exit-Self-Refresh Requirements Not Met**

### **Description**

When sequencing registered DIMMs out of self refresh state at the completion of an S1, S3 or LDTSTOP\_L initiated HyperTransport link width/frequency change, certain sequencing requirements of the registered DIMMs are not met.

### **Potential Effect on System**

Memory system failure leading to unpredictable system failure.

### **Suggested Workaround**

Do not use S1 S3 on a platform that employs registered DIMMs.

HyperTransport link width/frequency changes must be initiated using warm reset (as opposed to LDTSTOP\_L).

### **Fix Planned**

Yes

## **75 APIC Timer Accuracy Across Power Management Events**

### **Description**

The APIC timer may be inaccurate by up to 1  $\mu$ s across each use of S1 or LDTSTOP\_L initiated HyperTransport link width/frequency changes.

### **Potential Effect on System**

No observable system impact expected.

### **Suggested Workaround**

None.

### **Fix Planned**

No

## **76 APIC Timer Undercounts In Divide-by-8 Low Power Mode**

### **Description**

If S1 or LDTSTOP\_L initiated HyperTransport link width/frequency changes are performed with the Clock Divisor Select (ClkSel) set to divide-by-8, then the APIC timer incorrectly counts at 1/8 its intended rate. This miscounting remains in effect for as long as the processor remains in the divide-by-8 mode.

### **Potential Effect on System**

For S1, the divide-by-8 mode is not typically used, so no system implication is expected.

LDTSTOP\_L initiated HyperTransport link width/frequency changes do typically use the divide-by-8 ClkSel setting (to minimize latency) and are therefore affected by this erratum. However, for these operations the time spent in the divide-by-8 mode is limited to approximately 1  $\mu$ s per use, implying the APIC timer may lose approximately 0.875  $\mu$ s each time one of these transitions is performed. This error would be in addition to any other APIC timer accuracy errors that may exist.

### **Suggested Workaround**

None required. The accuracy loss is small and no observable system impact is expected.

### **Fix Planned**

Yes

## **77 Long Mode CALLF or JMPF May Fail To Signal GP When Callgate Descriptor is Beyond GDT/LDT Limit**

### **Description**

If the target selector of a far call or far jump (CALLF or JMPF) instruction references a 16-byte long mode system descriptor where any of the last 8 bytes are beyond the GDT or LDT limit, the processor fails to report a General Protection fault.

### **Potential Effect on System**

None expected, since the operating system typically aligns the GDT/LDT limit such that all descriptors are legal. However, in the case of erroneous operating system code, the above described GP fault will not be signaled, resulting in unpredictable system failure.

### **Suggested Workaround**

None required, it is anticipated that long mode operating system code will ensure the GDT and LDT limits are set high enough to cover the larger (16-byte) long mode system descriptors.

### **Fix Planned**

No

## **78 APIC Interrupt Latency With C2 Enabled**

### **Description**

If an APIC interrupt is delivered to the processor at a time when interrupts are masked (i.e., EFLAGS.IF=0), and just shortly before entering the C2 power state, then the interrupt may experience a long latency before being serviced.

The interrupt is not lost, but it is not serviced until some other wakeup event (for example, a timer tick) occurs to take the processor out of the C2 state.

### **Potential Effect on System**

Excessively long interrupt latencies may occur, resulting in unpredictable system failures.

### **Suggested Workaround**

Do not enable the C2 power state.

### **Fix Planned**

Yes

## **79 Power Management Limitations Above 1.50V**

### **Description**

Processor versions with a core voltage greater than 1.50V do not support Northbridge low power mode while in the S1 power state or LDTSTOP\_L initiated HyperTransport link width/frequency changes.

### **Potential Effect on System**

Unpredictable system failures may occur.

### **Suggested Workaround**

For affected versions of the processor, do not enable Northbridge low power mode in the S1 power state. Specifically, clear the NBLowPwrEn bit in SMAF code 011 of the Power Management Control Registers (i.e., clear Dev:3x80[25]).

Also, use warm reset (rather than LDTSTOP\_L) to initiate HyperTransport link width/frequency changes.

### **Fix Planned**

Yes

## **80 Registered DIMM Initialization Requirements Not Met**

### **Description**

When initializing registered DIMMs after a powerup or warm reset assertion, the time interval between the deassertion of MEMRESET\_L and the assertion of CKE is not sufficient for some DIMMs.

### **Potential Effect on System**

The memory system may fail to initialize, leading to boot failure.

### **Suggested Workaround**

A board level workaround is available for this problem, see *Methodologies for Using Registered DIMMs with AMD Athlon™ 64 and AMD Opteron™ Processors*, order #27510, for details.

### **Fix Planned**

Yes



## 81 Cache Coherency Problem with Hardware Prefetching and Streaming Stores

### Description

If the processor's hardware prefetch mechanism initiates a cache line prefetch at approximately the same time as a streaming store (MOVNT\* or MASKMOV\*) is performed to that same address, then a stale copy of that line may be loaded into the cache.

### Potential Effect on System

Cache coherency failure leading to unpredictable system failure. This erratum affects both uniprocessor and multiprocessor configurations. It has only ever been observed in a randomized diagnostic environment.

### Suggested Workaround

BIOS should set DC\_CFG.DIS\_SMC\_CHK\_BUF (bit 10 of MSR 0xC001\_1022) to disable the SMC check buffer for streaming stores. No loss of functionality occurs as a result of setting this bit.

### Fix Planned

Yes

## 82 Certain Faults On Far Transfer Instructions In 64-Bit Mode Save Incorrect RIP

### Description

This erratum affects the far transfer instructions (CALLF, RETF, IRET, JMPF) in 64-bit mode. If a far transfer is executed in 64-bit mode and:

- The RIP of the far transfer is 4 GB or greater (> 32 bits)
- The target is a 32-bit compatibility segment
- The far transfer encounters a fault *\_after\_* loading the CS

then the RIP pushed onto the exception handler stack will be erroneously truncated to 32-bits. The following table lists the instructions and faults that are subject to this erratum.

Instruction	Fault
CALLF (intersegment, no gate)	Target limit violation Faults on stack pushes
JMPF (intersegment, no gate)	Target limit violation
RETF/IRET (no CPL change)	Target limit violation
RETF/IRET (with CPL change)	Target limit violation Fault while loading new SS

### Potential Effect on System

The fault handler will return to the incorrect address if it attempts to IRET back to the faulting instruction (the far transfer), leading to unpredictable system failures.

### Suggested Workaround

None required. This erratum can only affect 64-bit operating systems, but has never been seen in such systems. It has only been seen in random instruction testing.

This erratum can only affect kernel operating system code, not applications since transfers from 64-bit mode to 32-bit compatibility can only be done by the kernel. Furthermore, existing 64-bit operating systems provide enough error checking and are constructed such that it is not possible to encounter one of the above faults at the point the kernel uses the far transfers to exit to a 32-bit compatibility application.

### Fix Planned

Yes

## 83 DC Machine Check Extended Error Code Bit Not Writeable

### Description

The extended error code bit in the DC Machine Check Status Register (i.e., bit 16 of MSR 0x0401) cannot be predictably updated by software. When this register is written, bit 16 may be updated to either a zero or a one depending on internal processor conditions.

The ability of this register to correctly log and classify machine check errors is not compromised by this erratum. When an error occurs, all status information (including the extended error code bit) is captured correctly and is readable by software. The erratum only affects the software writeability of this bit.

### Potential Effect on System

None expected. As described, errors are captured and classified correctly and are software readable. Since software only interprets the extended error code bit in the context of a valid TLB error, the inability of software to clear that extended error bit is of no functional consequence.

### Suggested Workaround

None required.

### Fix Planned

Yes

## **85 Insufficient Delay Between MEMCLK Startup and CKE Assertion During Resume From S3**

### **Description**

When sequencing the DRAMs out of self refresh during a resume from the S3 (Suspend-to-RAM) state, the processor fails to insert sufficient delay between MEMCLK startup and CKE assertion.

### **Potential Effect on System**

Memory system failure leading to unpredictable system failure.

### **Suggested Workaround**

Modify the resume from S3 BIOS sequence such that sufficient delay is inserted between the time MEMCLK is enabled in the DRAM Config High Register (Dev:2x94) and the time the ESR and SR\_S bits are written in the DRAM Config Low Register (Dev:2x90[13:12]) to exit self refresh. For registered DIMMs, 100  $\mu$ s of delay is required. For unbuffered DIMMs, 10  $\mu$ s of delay is required.

### **Fix Planned**

No

## **86 DRAM Data Masking Feature Can Cause ECC Failures**

### **Description**

Under certain conditions, the memory controller fails to generate a DRAM read request when performing partial writes to an already allocated write combining buffer. Because the DRAM is not read for these subsequent write requests, the generated ECC bits are incorrect.

### **Potential Effect on System**

Incorrect data can be read back from DRAM.

### **Suggested Workaround**

BIOS should disable the data masking feature when ECC DIMMs are used by setting the DisDatMsk bit (Northbridge Configuration Register - MSR C001\_001F, bit 36).

### **Fix Planned**

Yes

## 88 SWAPGS May Fail To Read Correct GS Base

### Description

The SWAPGS instruction fails to cause an input dependency on the GS segment register. If the GS segment register has been recently changed via a MOV or POP segment register instruction, SWAPGS may incorrectly save the old value of GS base into the KernelGSbase MSR.

### Potential Effect on System

The KernelGSbase MSR may be corrupted.

### Suggested Workaround

Between a MOV/POP into GS and any subsequent SWAPGS there must be a synchronizing operation. That operation can be one of the following:

- Any of the instructions that are required by the x86 architecture to be serializing (see *AMD64 Architecture Programmer's Manual Volume 2: System Programming*, order# 24593).
- A trap, interrupt or exception.
- An SFENCE or MFENCE instruction.
- An instruction that flushes the pipeline:
  - CALLF, JMPF, RETF, INTn, IRET, SYSCALL, SYSRET.

The MFENCE alternative is the lowest latency and the recommended alternative. The others are mentioned in case the code already satisfies them by construction.

### Fix Planned

Yes

## 89 Potential Deadlock With Locked Transactions

### Description

Downstream non-posted requests to devices that are dependent on the completion of an upstream non-posted request can cause a deadlock in the presence of transactions resulting in bus locks, as shown in the following two scenarios:

1. A downstream non-posted read to the LPC bus occurs while an LPC bus DMA is in progress. The legacy LPC DMA blocks downstream traffic until it completes its upstream reads.
2. A downstream non-posted read is sent to a device that must first send an upstream non-posted read before it can complete the downstream read.

In both cases, a locked transaction causes the upstream channel to be blocked, causing the deadlock condition.

### Potential Effect on System

The system fails due to a bus deadlock.

### Suggested Workaround

BIOS should set the DisIOReqLock bit (bit 3 in NB\_CFG, MSR C001\_001F).

### Fix Planned

No

## **90 False IC Machine Check Overflow Error Logged On Reset**

### **Description**

If a processor cold or warm reset occurs during a precise window when the instruction cache is being accessed due to a branch re-direct, a false IC Machine Check Overflow error may be logged in the IC Machine Check Status register.

### **Potential Effect on System**

System operation is not compromised, but a false machine check overflow error (bit 62) may be reported in the IC Machine Check Status register.

### **Suggested Workaround**

After a cold or warm reset, BIOS should clear the IC Machine Check Status register if the valid bit in that register is 0.

### **Fix Planned**

Yes



## 91 Software Prefetches May Report A Page Fault

### Description

Software prefetch instructions are defined to ignore page faults. Under highly specific and detailed internal circumstances, a prefetch instruction may report a page fault if both of the following conditions are true:

- The target address of the prefetch would cause a page fault if the address was accessed by an actual memory load or store instruction under the current privilege mode;
- The prefetch instruction is followed in execution-order by an actual or speculative byte-sized memory access of the same modify-intent to the same address.

PREFETCH and PREFETCHNTA/0/1/2 have the same modify-intent as a memory load access. PREFETCHW has the same modify-intent as a memory store access.

The page fault exception error code bits for the faulting prefetch will be identical to that for a byte-sized memory access of the same-modify intent to the same address.

Note that some misaligned accesses can be broken up by the processor into multiple accesses where at least one of the accesses is a byte-sized access.

If the target address of the subsequent memory access of the same modify-intent is aligned and not byte-sized, this errata does not occur and no workaround is needed.

### Potential Effect on System

An unexpected page fault may occur infrequently on a prefetch instruction.

### Suggested Workaround

Two workarounds are described for this erratum.

#### Kernel Workaround

The Operating System kernel can work around the erratum by allowing the page fault handler to satisfy the page fault to an "accessible" page regardless of whether the fault was due to a load, store, or prefetch operation. If the faulting instruction is permitted access to the page, return to it as usual. (An "accessible" page is one for which memory accesses are allowed under the current privilege mode once the page is resident in memory).

If the faulting instruction is trying to access an "inaccessible" page, scan the instruction stream bytes at the faulting Instruction Pointer to determine if the instruction is a prefetch. (An "inaccessible" page is one for which memory accesses are not allowed under the current privilege mode.) If the faulting instruction is a prefetch instruction, simply return back to it; the internal hardware conditions that caused the prefetch to fault should be removed and operation should continue normally. If it is not a prefetch instruction, generate the appropriate memory access control violation as appropriate. The performance impact of doing the scan is small because the actual errata is infrequent and does not produce an excessive number of page faults that affect system performance.

#### General Workaround

If the page-fault handler for a kernel can be patched as described in the preceding kernel workaround, no further action by software is required. The following general workarounds should only be considered for kernels where the page-fault handler can not be patched and a prefetch instruction could end up targeting an address in an "inaccessible" page.

Because the actual errata is infrequent, it does not produce an excessive number of page faults that affect system performance. Therefore a page fault from a prefetch instruction for an address within an "accessible" page does not require any general workaround.

Software can minimize the occurrence of the errata by issuing only one prefetch instruction per cache-line (a naturally-aligned 64-byte quantity) and ensuring one of the following:

- In many cases, if a particular target address of a prefetch is known to encounter this errata, simply change the prefetch to target the next byte.
- Avoid prefetching inaccessible memory locations, when possible.
- In the general case, ensure that the address used by the prefetch is offset into the middle of an aligned quadword near the end of the cache-line. For example, if the address desired to be prefetched is "ADDR", use an offset of 0x33 to compute the address used by the actual prefetch instruction as: "(ADDR & ~0x3f) + 0x33".

**Fix Planned**

Yes

## 92 Deadlock In Multi-Processor Systems May Occur When Earlier Operations Prevent An Older Store From Writing Data

### Description

A system deadlock may occur in multi-processor systems under the following conditions:

1. Interrupts are disabled.
2. A store operation occurs to a cacheable memory type.
3. The store is retired but not yet written the data cache.
4. The store is followed by a persistent (infinite) stream of loads while some of the loads are misaligned.
5. The misaligned loads are to the same cache index as the store (i.e., bits 11:6 are the same).
6. The misaligned loads are continually picked in the cycle preceding the store.
7. The destination cache line of the store is in a state other than modified (i.e., a probe from another processor to the same address as the store has previously transitioned this line to a shared state).

### Potential Effect on System

In the unlikely event that the above conditions occur, the system hangs.

### Suggested Workaround

None. This scenario was contrived in a highly randomized synthetic stress test and is not expected to occur in real systems.

### Fix Planned

Yes

## **93 RSM Auto-Halt Restart Returns To Incorrect RIP**

### **Description**

If an SMI occurs on a HALT instruction in 64-bit mode, and the subsequent RSM uses the Auto-Halt Restart feature, the RIP will incorrectly truncate to 32 bits.

### **Potential Effect on System**

Unpredictable system operation.

### **Suggested Workaround**

If the upper 32 bits of RIP are not cleared to zeroes, BIOS should clear the Auto-Halt Restart byte before executing the RSM instruction in the SMI handler. The RIP content can be verified and the Auto-Halt Restart byte can be modified in the SMM State-Save area in SM mode.

### **Fix Planned**

Yes

## 94 Sequential Prefetch Feature May Cause Incorrect Processor Operation

### Description

On an instruction cache miss, the sequential prefetch mechanism may enable the early prefetch of the next sequential cache line. Under a highly specific set of internal pipeline conditions this mechanism may cause the processor to hang or execute incorrect code in 64-bit systems running 32-bit compatibility mode applications.

### Potential Effect on System

Processor may deadlock or execute incorrect code.

### Suggested Workaround

BIOS should disable IC sequential prefetch for any system software which intends to operate in Long Mode, by setting IC\_CFG.DIS\_SEQ\_PREFETCH (bit 11 of MSR C001\_1021). System software should notify the BIOS what the expected operating mode is by using the Detect Target Operating Mode callback (INT 15, function EC00h) as described in the *BIOS and Kernel Developer's Guide for AMD Athlon™ 64 and AMD Opteron™ Processors*, order# 26094.

### Fix Planned

Yes

## **95 RET Instruction May Return To Incorrect EIP**

### **Description**

In order to efficiently predict return addresses, the processor implements a 12-deep return address stack to pair RETs with previous CALLs.

Under the following unusual and specific conditions, an overflowed hardware return stack may cause a RET instruction to return to an incorrect EIP in 64-bit systems running 32-bit compatibility mode applications:

- A CALL near is executed in 32-bit compatibility mode.
- Prior to a return from the called procedure, the processor is switched to 64-bit mode.
- In 64-bit mode, subsequent CALLs are nested 12 levels deep or greater.
- The lower 32 bits of the 64-bit return address for the 12th-most deeply nested CALL in 64-bit mode matches exactly the 32-bit return address of the original 32-bit mode CALL.
- A series of RETs is executed from the nested 64-bit procedures.
- The processor returns to 32-bit compatibility mode.
- A RET is executed from the originally called procedure.

### **Potential Effect on System**

The processor returns to an incorrect EIP, causing unpredictable system operation.

### **Suggested Workaround**

None. This has been observed in synthetic load stress-testing only, and not in any operating system or application.

### **Fix Planned**

Yes

## **96 Increased Memory Latency During P-State Changes**

### **Description**

The memory controller's idle counters are dynamically managed to help reduce page misses and conflicts. When LDTSTOP is asserted for a P-state (frequency) change, the memory controller incorrectly waits for the idle counters to expire before placing DRAM in self-refresh. This has the effect of increasing memory latency (up to 256 memory clocks in rare cases) during P-state changes.

### **Potential Effect on System**

The slight increase in memory latency may lead to performance anomalies depending on buffering capabilities of external devices.

### **Suggested Workaround**

None.

### **Fix Planned**

Yes

## **97 128-Bit Streaming Stores May Cause Coherency Failure**

### **Description**

Under a specific set of internal pipeline conditions, stale data may be left in the L1 cache when a 128-bit streaming store (MOVNT\*) to a writeback (WB) memory type misses in the L1 data cache and both L1 and L2 TLBs.

### **Potential Effect on System**

Memory coherence failures leading to unpredictable operation.

### **Suggested Workaround**

BIOS should set DC\_CFG.DIS\_CNV\_WC\_SSO (bit 3 of MSR 0xC001\_1022). The performance effects of setting this bit are limited to streaming stores to the write-combining (WC) memory type, a case expected to rarely occur in actual usage. No loss of performance occurs in the general case (WB memory type).

This workaround must not be applied to processors prior to revision C0.

### **Fix Planned**

Yes



## 98 LDTSTOP Assertion May Be Missed

### Description

If LDTSTOP width is too short relative to the programmed value of Clock Ramp Hysteresis, the LDTSTOP assertion may be missed.

### Potential Effect on System

FID changes or HyperTransport width/frequency changes may not work correctly.

### Suggested Workaround

Program the Clock Ramp Hysteresis value (Dev:3xD4 [10-8]) to be less than the LDTSTOP pulse width. Refer to the *BIOS and Kernel Developer's Guide for AMD Athlon™ 64 and AMD Opteron™ Processors*, order# 26094, for recommended LDTSTOP durations for each platform type.

### Fix Planned

Yes

## **99 Background Scrubbing Must Be Disabled With Non-Contiguous Memory Map**

### **Description**

On some systems, BIOS remaps part of system DRAM above the 4 Gbyte boundary for purposes of reclaiming memory lost to memory-mapped I/O resources. This results in a non-contiguous memory map. The background DRAM scrubber fails when it encounters an area of remapped memory, causing a system deadlock.

### **Potential Effect on System**

System deadlock when background scrubber is enabled.

### **Suggested Workaround**

Do not enable background scrubbing when the memory map is non-contiguous and a DIMM connected to chip select 0 is unpopulated.

### **Fix Planned**

Yes

## 100 Compatibility Mode Branches Transfer to Illegal Address

### Description

An IP-relative branch (call, jump, etc.), while in compatibility mode, may cause the processor to transfer execution to an illegal target address. A specific set of internal pipeline conditions and code alignment must be present. This transfer attempts to fetch an address outside the allowed compatibility mode 32-bit address range (addresses 0000\_0000h to FFFF\_FFFFh). The illegal target address will be within one of the following two 2GB logical address ranges:

0000\_0001\_0000\_0000h to 0000\_0001\_7FFF\_FFFFh

FFFF\_FFFF\_FFFF\_FFFFh to FFFF\_FFFF\_8000\_0000h

In most cases, due to the page mapping attributes that 64-bit operating systems place on these ranges, the branch immediately page faults. The page fault places the illegal address in CR2. However, the page fault also places a return instruction pointer on the exception stack that points to the correct target address.

This erratum does not occur in legacy mode nor in 64-bit long mode.

### Potential Effect on System

Compatibility mode code can transfer to an illegal address causing an unexpected page fault or incorrect results.

### Suggested Workaround

64-bit operating systems can work around this erratum by performing both of the following actions:

- Ensure that instruction fetches by compatibility user mode code will page fault when targeting the illegal logical address ranges shown above.
- From within the page fault handler, examine the page fault error code, the value of CR2, and the mode of the faulting instruction. If these indicate a compatibility mode instruction fetch to the above logical address ranges, then the handler can simply return. Because the faulting instruction pointer pushed on the stack by the processor is the correct target address, control will be transferred to that address and operation will continue normally. This workaround only needs to be applied to user mode page faults if the operating system allows compatibility mode only for user mode processes.

### Fix Planned

Yes

## **101 DRAM Scrubber May Cause Data Corruption When Using Node-Interleaved Memory**

### **Description**

When enabled, DRAM scrubbing cleans ECC errors from memory by performing a read-modify-write operation of each memory location at a programmable rate. In an MP system, if memory is configured as Node interleaved, the DRAM scrubber may cause data corruption. If another write request is made to the same address being accessed by the scrubber and occurs in a specific timing window, then the write data could be overwritten by older data from the scrubber.

### **Potential Effect on System**

Intermittent and non-repeatable data corruption may result from this erratum.

### **Suggested Workaround**

Disable either DRAM scrubbing or Node interleaving. BIOS should disable DRAM scrubbing (i.e., clear Dev:3x58[4:0] - DramScrub) if the Node interleaved memory option is selected. DRAM scrubbing and Node interleaving are setup options in the BIOS.

Note: The recommended BIOS default setting has been DRAM scrubber enabled and Node interleaving disabled.

### **Fix Planned**

Yes

## 103 AAM or DIV followed by AAM May Produce Incorrect Results

### Description

An AAM (ASCII Adjust after Multiply) or DIV instruction followed closely by an AAM instruction may interfere with each other. This erratum can only occur under a highly specific set of pipeline conditions when AAM is followed by another AAM within 3 instructions, or DIV is followed by AAM within 6 instructions. Either the first instruction (AAM or DIV) or the second instruction (AAM) may produce incorrect results.

### Potential Effect on System

Incorrect results may be produced. For the recommended usage of AAM, see section 3.3.3 of the *AMD64 Architecture Programmer's Manual Volume 1: Application Programming*, order# 24592.

### Suggested Workaround

This scenario was contrived in a highly randomized simulation environment and is not expected to occur in a real system. In the unlikely event that the erratum is observed, add NOP instructions to observe the instruction spacing given above.

### Fix Planned

Yes

## **104 DRAM Data Masking Feature Causes ChipKill ECC Failures When Enabled With x8/x16 DRAM Devices**

### **Description**

ChipKill ECC uses 16 ECC bits to correct and detect symbol errors within 128 data bits. If the data masking feature is enabled, the ECC bits are incorrectly calculated for partial writes to memory. This causes incorrect ECC to be written to DRAM. This erratum applies to x8 or greater DRAM devices, but not to x4 devices.

### **Potential Effect on System**

Data corruption may result from this erratum.

### **Suggested Workaround**

BIOS should disable the data masking feature when enabling ChipKill ECC with x8, x16 (or greater) ECC DIMMs, by setting the DisDatMsk bit (Northbridge Configuration Register - MSR C001\_001F, bit 36).

### **Fix Planned**

Yes

## 105 Misaligned 128-bit Store May Cause Deadlock

### Description

A processor deadlock may occur under the following conditions:

- A 128-bit store operation (MOVUPS, MOVUPD, MOVDQU) occurs to a cacheable memory type.
- The store is misaligned across two cache lines such that the upper 8 bytes span a cache line boundary.
- The store has retired but not yet written the data cache.
- The store is followed by two other load or store operations to the same cache index as the second half of the misaligned store (i.e., bits 11:6 are the same).

### Potential Effect on System

In the unlikely event that the above conditions occur the system may hang.

### Suggested Workaround

None.

### Fix Planned

Yes

## **106 Potential Deadlock with Tightly Coupled Semaphores in an MP System**

### **Description**

A write location may not become externally visible due to certain internal pipeline conditions involving tightly coupled semaphores across multiple processors:

1. Processor A does a write to clear processor B's semaphore but that write has not yet become visible to the system.
2. Processor B is waiting for its semaphore to be released before releasing processor A's semaphore.
3. Processor A immediately enters a spin loop waiting for its semaphore to be cleared by processor B, and the spin loop must fetch from the instruction cache (IC) on every cycle.
4. Because the IC is busy every cycle combined with other highly specific internal pipeline conditions, processor A's original write is prevented from being seen by processor B.

Additionally, event 3 (above) must follow event 1 closely in time and interrupts must be disabled.

### **Potential Effect on System**

The system may hang.

### **Suggested Workaround**

In an MP system, BIOS should set LS\_CFG.DIS\_LS2\_SQUISH (Bit 25 of MSR C001\_1020h).

### **Fix Planned**

Yes



## 107 Possible Multiprocessor Coherency Problem with Setting Page Table A/D Bits

### Description

In a multiprocessor system, a coherency failure may occur in a situation involving a TLB refill, an L1 fill, an L1 victim write, and an external probe, when all four addresses match as described in the following sequence:

1. A TLB miss occurs which requires the state of the page A (accessed) or D (Dirty) bit in one of the associated page map entries.
2. The cache line containing the page map entry must hit in the L2.
3. A younger load or store misses both in the data cache (DC) and L2 causing a DC line fill.
4. The DC fill generates an L1 victim.
5. The L1 victim is in the modified state.
6. The L1 fill matches the L1 index (11:6) of the TLB reload but not the L2 index [15:6, 14:6, or 13:6 for 1M, 512K, 256K] of the cache line containing the above page map entry.
7. The LRU bit for the L2 index points to the way (1 of 16) containing the page map entry.
8. The L1 victim arrives at the L2 in a small window after the TLB reload read, but before the write of the A/D bit(s).
9. An external probe arrives for the same address as the page map entry.

### Potential Effect on System

In the unlikely event that the above conditions occur, multiprocessor memory coherency issues may occur leading to unpredictable system failure. This erratum has not been observed outside of a highly randomized synthetic stress test.

### Suggested Workaround

BIOS should set BUCFG.ThrL2IdxCmpDis (bit 43) to one.

### Fix Planned

Yes

## **108 CPUID Instruction May Return Incorrect Model Number In Some Processors**

### **Description**

The Model Number field (bits 7-4 of the value returned in EAX by CPUID Function 1) may report an Eh for Revision DH7-CG processors in a 754-pin package. The correct Model Number value for this processor is Ch. The value read from the Model Number field may not be consistent and could report a value of Ch or Eh over time.

Note this affects the Model Number field only, all other CPUID fields are unaffected.

### **Potential Effect on System**

The effect of this erratum is similar to a new processor being installed in a system that has not been updated with BIOS and system software that recognizes the new model number:

- System software dependent on the Model Number field of the CPUID may fail to recognize the processor causing processor-specific drivers not to load.
- BIOS may fail to recognize the processor and take a safe default action.
- Application software typically does not require knowledge of the Model Number field and therefore is unaffected.

### **Suggested Workaround**

BIOS and system software should be modified to recognize the Model Number field values of both Ch and Eh as referring to Revision DH7-CG processors in a 754-pin package.

### **Fix Planned**

Yes

## 109 Certain Reverse REP MOVS May Produce Unpredictable Behavior

### Description

In certain situations a REP MOVS instruction may lead to incorrect results. An incorrect address size, data size or source operand segment may be used or a succeeding instruction may be skipped. This may occur under the following conditions:

- EFLAGS.DF=1 (the string is being moved in the reverse direction).
- The number of items being moved (RCX) is between 1 and 20.
- The REP MOVS instruction is preceded by some microcoded instruction that has not completely retired by the time the REP MOVS begins execution. The set of such instructions includes BOUND, CLI, LDS, LES, LFS, LGS, LSS, IDIV, and most microcoded x87 instructions.

### Potential Effect on System

Incorrect results may be produced or the system may hang.

### Suggested Workaround

Contact your AMD representative for information on a BIOS update.

### Fix Planned

Yes

## **111 Rtt Specification Violation**

### **Description**

The Rtt specification for the HyperTransport™ pins may be violated on some processor revisions.

### **Potential Effect on System**

Potential violations of the VID (input differential voltage) and Tr/Tf (slew rate) HyperTransport specifications. There are no known failures related to this problem.

### **Suggested Workaround**

None required.

### **Fix Planned**

Yes

## 112 Self-Modifying Code May Execute Stale Instructions

### Description

The processor may execute stale instructions in a situation involving a store instruction modifying a younger instruction within two cache lines of each other and an internal processor trap condition occurring in a small window:

1. The fetch for the store and the younger instruction must hit in the instruction cache (IC) and are brought into the processor pipeline.
2. The store speculatively executes, and prior to invalidating the line in the IC, an internal trap event occurs.
3. This internal trap event must redirect the instruction fetch to a point in the instruction stream just before the store.
4. The fetch associated with the redirect must hit in the IC before the first execution of the store (in step 2) has invalidated the line containing the target.
5. During the small window of time between the refetch of the younger instruction and the fetch of the last instruction in that cache line from the IC, the speculative store from step 2 invalidates the IC line which contains the younger instruction.

### Potential Effect on System

In the unlikely event that the above conditions occur, the processor will execute stale instruction(s).

Just in time (JIT) compilers lack the proximity of the store instruction to the modified code and thus are not affected.

### Suggested Workaround

This scenario was contrived in a highly randomized simulation environment and is not expected to occur in a real system. In the unlikely event that the erratum is observed, a serializing instruction can be inserted prior to executing the modified code.

### Fix Planned

Yes

## **117 Incorrect Value May Be Returned When Reading ChipKill ECC Syndrome**

### **Description**

The 16 bits of the ChipKill ECC syndrome are logged in the MC4\_STATUS register.

During a small window while the processor is updating the syndrome bits for a correctable or uncorrectable error, a read of the MC4\_STATUS register may return an incorrect syndrome value.

This erratum does not occur if the MC4\_STATUS register already contains information on an uncorrectable error.

### **Potential Effect on System**

ECC logging software may not be able to uniquely identify the DIMM with the error. However, the failing address that is logged allows the error to be isolated to a pair of DIMMs.

### **Suggested Workaround**

Machine check logging software can read the MC4\_STATUS register multiple times until a consistent value is obtained. Since the failing window is short and occurs only when encountering a Chipkill ECC error during the read, a consistent value is immediately obtained in nearly all cases.

### **Fix Planned**

Yes

## 121 Sequential Execution Across Non-Canonical Boundary Causes Processor Hang

### Description

The processor will hang when the following conditions are met:

- The processor is in 64-bit mode
- The code segment limit = 0xFFFF\_FFFF
- The last byte of the current instruction is located at 0x7FFF\_FFFF\_FFFF
- The next sequential instruction fetch is attempted at 0x8000\_0000\_0000

The correct behaviour is to cause #GP (general protection exception).

### Potential Effect on System

The system hangs.

### Suggested Workaround

The operating system should not allocate the page at the boundary of canonical address space (user).

### Fix Planned

Yes

## 122 TLB Flush Filter May Cause Coherency Problem in Multiprocessor Systems

### Description

Under highly specific internal timing conditions in a multiprocessor configuration, coherency problems may arise between the page tables in memory and the translations stored in the on-chip TLBs. This can result in the possible use of stale translations even after software has performed a TLB flush.

### Potential Effect on System

Unpredictable system failure. This scenario has only been observed in a highly randomized synthetic stress test.

### Suggested Workaround

In multiprocessor systems, disable the TLB flush filter by setting HWCR.FFDIS (bit 6 of MSR 0xC001\_0015).

### Fix Planned

Yes



## Documentation Support

---

The following documents provide additional information regarding the operation of the processor:

- *AMD Athlon™ 64 Product Data Sheet*, order# 24659
- *AMD Opteron™ Product Data Sheet*, order# 23932
- *Mobile AMD Athlon™ 64 Product Data Sheet*, order# 27105
- *Mobile AMD Athlon™ XP-M Processor Family 15 Product Data Sheet*, order# 31511
- *AMD Athlon™ 64 FX Product Data Sheet*, order# 30431
- *CPUID Guide for AMD Athlon™ 64 and AMD Opteron™ Processors: Application Note Addendum*, order# 25481
- *BIOS and Kernel Developer's Guide for AMD Athlon™ 64 and AMD Opteron™ Processors*, order# 26094
- *AMD Athlon™ 64 Processor Motherboard Design Guide*, order# 24665
- *AMD64 Architecture Programmer's Manual Volume 1: Application Programming*, order# 24592
- *AMD64 Architecture Programmer's Manual Volume 2: System Programming*, order# 24593
- *AMD64 Architecture Programmer's Manual Volume 3: Instruction-Set Reference*, order# 24594
- *AMD64 Architecture Programmer's Manual Volume 4: 128-Bit Media Instructions*, order# 26568
- *AMD64 Architecture Programmer's Manual Volume 5: 64-Bit Media and x87 Floating-Point Instructions*, order# 26569
- *Methodologies for Using Registered DIMMs with AMD Athlon™ 64 and AMD Opteron™ Processors*, order# 27510

See the AMD Web site at [www.amd.com](http://www.amd.com) for the latest updates to documents. For documents subject to a non-disclosure agreement (NDA), please contact your local sales representative.