



Revision Guide for AMD Opteron™ Processors

| | |
|--------------------------------|-----------------------|
| Publication # 25759 | Revision: 3.07 |
| Issue Date: August 2003 | |

Preliminary Information

© 2003 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, 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

| Date | Revision | Description |
|-------------|----------|--|
| August 2003 | 3.07 | Added erratum #95, #97, and #98; Modified Description in errata #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 Opteron™ Processors

The purpose of the *Revision Guide for AMD Opteron™ Processors* is to communicate updated product information to designers of computer systems and software developers. 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 AMD Opteron™ processor to deviate from the published specifications.
- **Documentation Support:** This section, starting on page 42, provides a listing of available technical support resources.

Revision Guide Policy

Occasionally, AMD identifies product errata that cause the AMD Opteron processor to deviate from published specifications. Descriptions of identified product errata are designed to assist system and software designers in using the AMD Opteron processor. Furthermore, 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 Processor

| Revision | CPUID Function 1 EAX Value AMD Opteron™ Processor |
|----------|--|
| SH7-B3 | 00000F51h |
| SH7-C0 | 00000F58h |

Product Errata

This section documents product errata for AMD Opteron™ 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. “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 |
| 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 |
| 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 RDPIC 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 | |
| 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 | |

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

| No. | Errata Description | Revision Number | |
|-----|--|-----------------|--------|
| | | SH7-B3 | SH7-C0 |
| 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 |
| 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 |
| 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 |
| 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 |
| 94 | Sequential Prefetch Feature May Cause Incorrect Processor Operation | X | X |
| 95 | RET Instruction May Return To Incorrect EIP | X | X |
| 97 | 128-Bit Streaming Stores May Cause Coherency Failure | | X |
| 98 | LDTSTOP Assertion May Be Missed | | 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

If the described software scenario can be avoided, this erratum will not occur (i.e., if software never maps the same linear page global in one context and not global in another context). If the software scenario cannot be avoided, two other potential workarounds are possible:

1. Disable TLB error reporting in the machine check architecture (or disable the machine check architecture entirely). 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.
2. Disable the TLB flush filter by setting HWCR.FFDIS (bit 6) in MSR C001_0015h.

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 or 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 μ 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 μ s per use, implying the APIC timer may lose approximately 0.875 μ 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 an 8-byte long mode system descriptor where any of the last 4 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 signalled, 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 (8-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

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

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 *AMD x86-64 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, the following conditions may cause the prefetch instruction to report a page fault when the address would normally cause a page fault:

- A PREFETCH or PREFETCHNTA/0/1/2 followed by a load to the same address.
- A PREFETCHW followed by a store to the same address.

Potential Effect on System

An unexpected page fault may occur on a prefetch instruction.

Suggested Workaround

Disable speculative TLB reloads by setting DC_CFG bit 4.

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

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 by setting IC_CFG.DIS_SEQ_PREFETCH (bit 11 of MSR C001_1021).

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

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

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

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 when HyperTransport link width/frequency changes are used by the system.

Fix Planned

Yes

Documentation Support

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

- *AMD Processor Recognition Application Note*, order# 20734
- *BIOS and Kernel Developer's Guide for AMD Athlon™ 64 and AMD Opteron™ Processors*, order# 26094
- *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 Opteron™ Processors*, order# 27510
- *AMD Opteron™ Processor Data Sheet*, order# 23932

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