



AMD Athlon™ Processor Model 6 Revision Guide

Publication # 24332 Rev: G
Issue Date: **October 2003**

Preliminary Information

© 2001–2003 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.

The AMD products described herein may contain design defects or errors (“Product Errata”) that cause AMD products to deviate from published specifications. Currently characterized Product Errata may be available upon request.

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, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

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

Revision History

Date	Rev	Description
October 2003	G	Revised erratum #25.
June 2003	F	Added erratum #25.
December 2002	E	Added errata #22–24.
July 2002	D	Added errata #20 and #21.
October 2001	C	<ul style="list-style-type: none"> ■ Added silicon revision A5 information ■ Added erratum #18 and #19 ■ Added Table 2: Cross-reference of Erratum to Processor Segments to include all processor segments ■ Updated reference documents listed in section 3.1
May 2001	B	Initial public release.

AMD Athlon™ Processor Model 6 Revision Guide

The purpose of the *AMD Athlon™ Processor Model 6 Revision Guide* is to communicate updated product information on the AMD Athlon processor model 6 to designers of computer systems and software developers. This revision guide applies to the AMD Athlon processor model 6, mobile AMD Athlon processor model 6, and AMD Athlon MP processor model 6. This guide consists of three sections:

- **Product Errata:** This section, which starts on page 5, 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. A product errata may cause the behavior of the AMD Athlon processor model 6 to deviate from the published specifications.
- **Revision Determination:** This section, which starts on page 18, shows the AMD Athlon processor model 6 identification numbers returned by the CPUID instruction for each revision of the processor.
- **Technical and Documentation Support:** This section, which starts on page 19, provides a listing of available technical support resources. It also lists corrections, modifications, and clarifications to listed documents. Please refer to the data sheets listed in this section for product marking information.

Revision Guide Policy

Occasionally AMD identifies deviations from or changes to the specification of the AMD Athlon processor model 6. These changes are documented in the *AMD Athlon™ Processor Model 6 Revision Guide* as errata. Descriptions are written to assist system and software designers in using the AMD Athlon processor model 6 and corrections to AMD's documentation on the AMD Athlon processor model 6 are included. This release documents currently characterized product errata.

1 Product Errata

This section documents AMD Athlon™ Processor Model 6 product errata. The errata are divided into categories to assist referencing particular errata. 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 1 cross-references the revisions of the processor to each erratum. An “X” indicates that the erratum applies to the stepping. The absence of an “X” indicates that the erratum does not apply to the stepping. Table 2 on page 6 cross-references erratum to each processor segment. An “X” indicates that the erratum applies to the processor segment.

Note: *There can 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 1. Cross-Reference of Product Revision to Errata

Errata Numbers and Description	Revision Numbers		
	A0	A2	A5
16 INVLPG Instruction Does Not Flush Entire Four-Megabyte Page Properly with Certain Linear Addresses	X	X	
17 Deadlock May Occur in a Two-Processor System in the Presence of Probe to Memory- Mapped I/O	X	X	X
18 Processor May Issue Non-Connect Bus Cycle After FID Special Cycle	X	X	
19 Processor Does Not Support Reliable Microcode Patch Mechanism			X
20 Processor Performance Counters Do Not Count Some x86 Instructions	X	X	X
21 A Speculative SMC Store Followed by an Actual SMC Store May Cause One-Time Stale Execution	X	X	X
22 Real Mode RDPMC with Illegal ECX May Cause Unpredictable Operation	X	X	X
23 Using Task Gates With Breakpoints Enabled May Cause Unexpected Faults	X	X	X
24 Single Step Across I/O SMI Skips One Debug Trap	X	X	X
25 Software Prefetches May Report A Page Fault	X	X	X

Table 2. Cross-Reference of Erratum to Processor Segments

Errata Number	Workstation/Server¹	Desktop²	Mobile³
16	X	X	X
17	X		
18			X
19	X	X	X
20	X	X	X
21	X	X	X
22	X	X	X
23	X	X	X
24	X	X	X
25	X	X	X

Notes:

1. The workstation/server segment currently includes the AMD Athlon™ MP processor.
2. The desktop segment currently includes the AMD Athlon XP processor.
3. The mobile segment currently includes the AMD Athlon 4 processor.

16 INVLPG Instruction Does Not Flush Entire Four-Megabyte Page Properly with Certain Linear Addresses

Products Affected. A0, A2

Normal Specified Operation. After executing an INVLPG instruction the TLB should not contain any translations for any part of the page frame associated with the designated logical address.

Non-conformance. When the logical address designated by the INVLPG instruction is mapped by a 4-Mbyte page mapping and LA[21] is equal to one it is possible that the TLB will still retain translations after the instruction has finished executing.

Potential Effect on System. The residual data in the TLB can result in unexpected data access to stale or invalid pages of memory.

Suggested Workaround. When using the INVLPG instruction in association with a page that is mapped via a 4-Mbyte page translation, always clear bit 21.

Resolution Status. Fix planned for a future revision.

17 Deadlock May Occur in a Two-Processor System in the Presence of Probe to Memory-Mapped I/O

Products Affected. A0, A2, A5

Normal Specified Operation. Processor should not hang.

Non-conformance. In a multiprocessor system, if one processor (A) is continuously writing to a cacheable memory-mapped I/O block while the other processor (B) is trying to read the same cacheable I/O block, and at the same time both processors are also trying to write a different memory-based cache block, then processor B may hang. Should this occur and processor A fields an interrupt, the deadlock will be resolved.

Potential Effect on System. System will hang or exhibit performance degradation.

Suggested Workaround. The current processor design assumes that memory-mapped I/O is incoherent and does not handle all deadlock cases. System logic should not generate probes for memory mapped I/O addresses.

Resolution Status. No fix planned.

18 Processor May Issue Non-Connect Bus Cycle After FID Special Cycle

Products Affected. A0, A2

Normal Specified Operation. The first processor cycle after a FID Change special cycle should be a Connect special cycle.

Non-conformance. In rare circumstances, a processor victim write may be pending inside the processor when the FID Change special cycle is issued. Several bus clocks later, the WrVictimBlk command for the victim will be issued. This violates the specification, which states that all processor-based commands should be finished before the FID change special cycle.

Potential Effect on System. The core logic may become confused.

Suggested Workaround. System core logic can wait for numerous bus clocks after receiving the FID change special cycle before attempting to disconnect in order to generate a window sufficiently large enough to allow the WrVictimBlk transaction to take place prior to the disconnect.

Resolution Status. Fix planned for a future revision.

19 Processor Does Not Support Reliable Microcode Patch Mechanism

Products Affected. A5

Normal Specified Operation. The processor should function properly after a microcode patch is loaded.

Non-conformance. The processor has the patch RAM BIST function disabled. Since BIST is not run on the patch RAM, reliable operation of the patch RAM cannot be guaranteed. Therefore it should not be used.

Potential Effect on System. When a microcode patch is loaded, the system may not behave properly.

Suggested Workaround. Do not load a microcode patch.

Resolution Status. Fix planned for a future revision.

20 Processor Performance Counters Do Not Count Some x86 Instructions

Products Affected. A0, A2, A5

Normal Specified Operation. The processor should count all x86 instructions when programmed to do so.

Non-conformance. There are two types of uncounted instructions. One set of instructions is always uncounted. Another set of instructions are uncounted only if a certain data dependency exists.

Instructions never counted are: RDMSR, WRMSR, FSTENV, FSAVE, FLDENV, FPTAN, FYL2XP1, FCLEX, LLDT, LTR, MOV CR_x, LGDT, LIDT, INVLPG, INVD, WBINVD, MOV DR_x, CPUID, and SFENCE.

Instructions that are uncounted only when certain data dependencies exist are:

- LAR, LSL, VERR, VERW if they clear the Zero Flag
- FXSAVE, FXRSTOR if FERR is changed
- FPU instructions with exceptional data conditions
- IO instructions that detect an interrupt
- POPF with the trap flag =1
- POPFD and PUSHFD with IOPL not equal 3 and Virtual Mode enabled
- POPFD when Alignment Check is being enabled
- MOV SS with the trap flag =1
- Segment Loads that generate accessed bit exceptions
- STI with the trap flag or the interrupt flag already a 1
- CLTS with the CR0.TS flag =1
- LMSW that changes any bit

Potential Effect on System. Performance counter may under count the actual number of x86 instructions.

Suggested Workaround. Versions of the AMD Athlon™ processor not affected by this erratum may be used to gather instruction counts.

Resolution Status. No fix planned.

21 A Speculative SMC Store Followed by an Actual SMC Store May Cause One-Time Stale Execution

Products Affected. A0, A2, A5

Normal Specified Operation. Self-modifying code sequences should be correctly detected and handled in a manner consistent with canonical results; stale code should not be executed.

Non-conformance. The following scenario can result in a one-time execution of stale instructions:

1. A speculative store instruction initiates a request (R) to modify a 64-byte cache line with address A, which currently resides within the L1 instruction cache.
2. The speculative store instruction is ultimately not executed because of a branch misprediction. However, the store R is still in flight attempting to bring the line into the data cache in the modified state.
3. The instruction cache, which fetches instructions 16 bytes at a time, is redirected by the branch into the cache line with address A and fetches a portion of the line into the instruction buffer.
4. R then invalidates the instruction cache line with address A and brings the line into the L1 data cache, marking it as modified. However, the instruction buffer, which also contains some bytes from address A, is not invalidated.
5. The instruction fetch mechanism attempts to read the next 16-byte chunk of code and must issue a request to bring the 64-byte line back into the instruction cache.
6. This instruction cache request for address A hits on the modified line now in the L1 cache, and evicts it from the data cache to the L2.
7. A second store instruction (S) from the instruction buffer is issued into the execution units. S is a self-modifying code reference to another instruction that currently exists in the 64-byte cache block at address A and is also in the instruction buffer.
8. The execution of S detects that an instruction request to fetch address A is in flight. However, the store request is given priority. Since it now hits in the L2 and the L2 state is modified, it assumes that the line cannot be in the instruction cache or the instruction buffer.

Potential Effect on System. The processor will execute stale code instructions.

Suggested Workaround. None. This failure has only been observed in internally generated synthetic code.

Resolution Status. No fix planned.

22 Real Mode RDPMC with Illegal ECX May Cause Unpredictable Operation

Products Affected. A0, A2, A5

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

Non-conformance. If the RDPMC is executed in real mode with a specific illegal value of $ECX=4$, then the processor may incorrectly enter the GP fault handler as if it were in 32-bit real 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 legal counter values (0-3). This circumstance is not expected to occur in normal operation and has only been detected in a simulation environment.

Resolution Status. No fix planned.

23 Using Task Gates With Breakpoints Enabled May Cause Unexpected Faults

Products Affected. A0, A2, A5

Normal Specified Operation. Task gates should correctly use the TSS selector out of the task gate for CALL and JMP instructions.

Non-conformance. 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, under certain timing scenarios, incorrectly use the new TSS base[15:0] contained in the new TSS as a selector. This will most likely cause a GP fault with an error code of the new TSS base.

Potential Effect on System. System failure.

Suggested Workaround. When using software that uses task gates with CALL or JMP instructions, do not enable breakpoints.

Resolution Status. No fix planned.

24 Single Step Across I/O SMI Skips One Debug Trap

Products Affected. A0, A2, A5

Normal Specified Operation. 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.

Non-conformance. 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.

Resolution Status. No fix planned.

25 Software Prefetches May Report A Page Fault

Products Affected. A0, A2, A5

Normal Specified Operation. Software prefetches should not report page faults if they encounter them.

Non-conformance. 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".

Resolution Status. No fix planned.

2 Revision Determination

Table 3 shows the AMD Athlon™ processor model 6 identification number returned by the CPUID instruction for each revision of the processor.

Table 3. CPUID Values for the Revisions of the AMD Athlon™ Processor Model 6

Revision	CPUID
A0	660
A2	661
A5	662

3 Technical and Documentation Support

The following documents provide additional information regarding the operation of the AMD Athlon™ processor model 6. Please refer to the data sheets listed in this section for product marking information.

- *AMD Athlon™ XP Processor Data Sheet: Processor Model 6*, order #24309
- *Mobile AMD Athlon™ 4 Processor Model 6 CPGA Data Sheet*, order #24319
- *AMD Athlon™ MP Processor Model 6 Data Sheet Multiprocessor-Capable for Workstation and Server Platforms*, order #24685
- *AMD Athlon™ MP Processor OPGA Data Sheet for Multiprocessor Platforms: Processor Model 6*, order #25480

For the latest updates, refer to www.amd.com and download the appropriate files. For documents under NDA, please contact your local sales representative for updates.