



Revision Guide for AMD Family 19h Models 00h-0Fh Processors

Publication # 56683	Revision: 1.08
Issue Date: August 2023	

Advanced Micro Devices 

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

Trademarks

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

Linux is a registered trademark of Linus Torvalds.

Microsoft and Windows are registered trademarks of Microsoft Corporation.

PCIe and PCI Express are registered trademarks of PCI-SIG.

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

List of Figures

Figure 1. Format of CPUID Fn0000_0001_EAX.....	8
--	---

List of Tables

Table 1. Arithmetic and Logic Operators.....	7
Table 2. CUID Values for AMD Family 19h Models 00h-0Fh SP3 Processor Revisions.....	8
Table 3. OSVW ID Length Register.....	9
Table 4. OSVW Status Register.....	9
Table 5. Cross Reference of Product Revision to OSVW ID.....	10
Table 6. Cross-Reference of Processor Revision to Errata.....	11
Table 7. Cross-Reference of Errata to Package Type.....	14
Table 8. Cross-Reference of Errata to Processor Segments.....	17

Revision History

Date	Revision	Description
August 2023	1.08	Added #1433, #1450, #1455, #1464, #1467, #1476, #1488, #1489
April 2022	1.07	Added #1285, #1363, #1365, #1373, #1378, #1379, #1380, #1381, #1386, #1390, #1391, #1394, #1407, #1415. Updated #1275.
July 2021	1.04	Initial public release.

Overview

The purpose of the *Revision Guide for AMD Family 19h Models 00h-0Fh* is to communicate updated product information to designers of computer systems and software developers. This revision guide includes information on the following products:

- AMD EPYC™ 7003 Series Processors
- AMD EPYC™ 7003 Series Processors with AMD 3D V-Cache

Feature support varies by brands and OPNs (Ordering Part Number). To determine the features supported by your processor, contact your customer representative.

This guide consists of these major sections:

- [Processor Identification](#) shows how to determine the processor revision and workaround requirements, and to construct, program, and display the processor name string.
- [Product Errata](#) 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](#) 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.

Conventions

Numbering

- **Binary numbers.** Binary numbers are indicated by appending a "b" at the end, e.g., 0110b.
- **Decimal numbers.** Unless specified otherwise, all numbers are decimal. This rule does not apply to the register mnemonics.
- **Hexadecimal numbers.** Hexadecimal numbers are indicated by appending an "h" to the end, e.g., 45F8h.
- **Underscores in numbers.** Underscores are used to break up numbers to make them more readable. They do not imply any operation. e.g., 0110_1100b.
- **Undefined digit.** An undefined digit, in any radix, is notated as a lower case "x".

Arithmetic and Logical Operators

In this document, formulas follow some Verilog conventions as shown in [Table 1](#).

Table 1. Arithmetic and Logic Operators

Operator	Definition
{ }	Curly brackets are used to indicate a group of bits that are concatenated together. Each set of bits is separated by a comma. E.g., {Addr[3:2], Xlate[3:0]} represents a 6-bit value; the two MSBs are Addr[3:2] and the four LSBs are Xlate[3:0].
	Bitwise OR operator. E.g. (01b 10b == 11b).
	Logical OR operator. E.g. (01b 10b == 1b); logical treats multibit operand as 1 if >=1 and produces a 1-bit result.
&	Bitwise AND operator. E.g. (01b & 10b == 00b).
&&	Logical AND operator. E.g. (01b && 10b == 1b); logical treats multibit operand as 1 if >=1 and produces a 1-bit result.
^	Bitwise exclusive-OR operator; sometimes used as "raised to the power of" as well, as indicated by the context in which it is used. E.g. (01b ^ 10b == 11b). E.g. (2^2 == 4).
~	Bitwise NOT operator (also known as one's complement). E.g. (~10b == 01b).
!	Logical NOT operator. E.g. (!10b == 0b); logical treats multibit operand as 1 if >=1 and produces a 1-bit result.
==	Logical "is equal to" operator.
!=	Logical "is not equal to" operator.
<=	Less than or equal operator.
>=	Greater than or equal operator.
*	Arithmetic multiplication operator.
/	Arithmetic division operator.
<<	Shift left first operand by the number of bits specified by the 2nd operand. E.g. (01b << 01b == 10b).
>>	Shift right first operand by the number of bits specified by the 2nd operand. E.g. (10b >> 01b == 01b).

Register References and Mnemonics

In order to define errata workarounds it is sometimes necessary to reference processor registers. References to registers in this document use a mnemonic notation consistent with that defined in the *Processor Programming Reference (PPR) for AMD Family 19h Model 00h-0Fh Processors*, order# 55898.

Processor Identification

This section shows how to determine the processor revision.

Revision Determination

A processor revision is identified using a unique value that is returned in the EAX register after executing the CPUID instruction function 0000_0001h (CPUID Fn0000_0001_EAX). Figure 1 shows the format of the value from CPUID Fn0000_0001_EAX.




Figure 1. Format of CPUID Fn0000_0001_EAX

The following tables show the identification numbers from CPUID Fn0000_0001_EAX for each revision of the processor to each processor segment. "X" signifies that the revision has been used in the processor segment. "N/A" signifies that the revision has not been used in the processor segment.

Table 2. CPUID Values for AMD Family 19h Models 00h-0Fh SP3 Processor Revisions

CPUID Fn0000_0001_EAX	AMD EPYC™ 7003 Series Processors	AMD EPYC™ 7003 Series Processors with AMD 3D V-Cache
00A00F11h (Milan-B1)	X	
00A00F12h (Milan-B2)		X

Mixed Processor Revision Support

AMD Family 19h processors with different revisions may not be mixed in a multiprocessor system.

Programming and Displaying the Processor Name String

This section, intended for system software programmers, describes how to program and display the 48-character processor name string that is returned by CPUID Fn8000_000[4:2]. The hardware or cold reset value of the processor name string is 48 ASCII NUL characters, so system software must program the processor name string before any general purpose application or operating system software uses the extended functions that read the name string. It is common practice for system software to display the processor name string and model number whenever it displays processor information during boot up.

Note: Motherboards that do not program the proper processor name string and model number will not pass AMD validation and will not be posted on the AMD Recommended Motherboard Web site.

The name string must be ASCII NUL terminated and the 48-character maximum includes that NUL character.

The processor name string is programmed by MSR writes to the six MSR addresses covered by the range MSRC001_00[35:30]h. Refer to the PPR for the format of how the 48-character processor name string maps to the 48 bytes contained in the six 64-bit registers of MSRC001_00[35:30].

The processor name string is read by CPUID reads to a range of CPUID functions covered by CPUID Fn8000_000[4:2]. Refer to CPUID Fn8000_000[4:2] in the PPR for the 48-character processor name string mapping to the 48 bytes contained in the twelve 32-bit registers of CPUID Fn8000_000[4:2].

Operating System Visible Workarounds

This section describes how to identify operating system visible workarounds.

MSRC001_0140 OS Visible Work-around MSR0 (OSVW_ID_Length)

This register, as defined in *AMD64 Architecture Programmer's Manual Volume 2: System Programming*, order# 24593, is used to specify the number of valid status bits within the OS Visible Work-around status registers.

The reset default value of this register is 0000_0000_0000_0000h.

System software shall program the OSVW_ID_Length to 0005h prior to hand-off to the OS.

Table 3. OSVW ID Length Register

Bits	Description
63:16	Reserved.
15:0	OSVW_ID_Length: OS visible work-around ID length. Read-write.

MSRC001_0141 OS Visible Work-around MSR1 (OSVW_Status)

This register, as defined in *AMD64 Architecture Programmer's Manual Volume 2: System Programming*, order# 24593, provides the status of the known OS visible errata. Known errata are assigned an OSVW_ID corresponding to the bit position within the valid status field.

Operating system software should use MSRC001_0140 to determine the valid length of the bit status field. For all valid status bits: 1=Hardware contains the erratum, and an OS software work-around is required or may be applied instead of a system software workaround. 0=Hardware has corrected the erratum, so an OS software work-around is not necessary.

The reset default value of this register is 0000_0000_0000_0000h.

Table 4. OSVW Status Register

Bits	Description
63:5	OsvwStatusBits: Reserved. OS visible work-around status bits. Read-write.
4	OsvwId4: Reserved, must be zero.
3	OsvwId3: Reserved, must be zero.
2	OsvwId2: Reserved, must be zero.
1	OsvwId1: Reserved, must be zero.

Table 4. OSVW Status Register (continued)

Bits	Description
0	OsvwId0: Reserved, must be zero.

System software shall program the state of the valid status bits as shown in [Table 5](#) prior to hand-off to the OS.

Table 5. Cross Reference of Product Revision to OSVW ID

CPUID Fn0000_0001_EAX (Mnemonic)	MSRC001_0141 Bits
00A00F11h (Milan-B1)	0000_0000_0000_0000h
00A00F12h (Milan-B2)	0000_0000_0000_0000h

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. This table cross-references the revisions of the part to each erratum. "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 do not affect this product family do not appear. In addition, 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 6. Cross-Reference of Processor Revision to Errata

No.	Errata Description	CPUTID Fn0000_0001_EAX	
		00A00F11h (Milan-B1)	00A00F12h (Milan-B2)
1155	DMA or Peer-to-peer Accesses Using Guest Physical Addresses (GPAs) May Cause IOMMU Target Abort	No fix planned	
1159	Writes to Base Frequency Register May be Ignored	No fix planned	
1160	SdpParity and XiVictimQueue Mask Bits Incorrectly Mask Additional Errors	No fix planned	
1169	PCIe® Error Masking May Fail to Mask Errors	No fix planned	
1193	Page Remapping Without Invalidation May Cause Missed Detection of Self-Modifying Code	No fix planned	
1197	IBS (Instruction Based Sampling) Register State May be Incorrect After Restore From CC6	No fix planned	
1200	xHCI Host May Hang If Full Speed or High Speed USB Hub is Connected	No fix planned	
1211	PCIe Incorrectly Logs TLP Prefix AER (Advanced Error Reporting) Register After Egress Blocking Error	No fix planned	
1216	IOMMU May Not Re-Walk Page Tables on a Present/Permission Fault	No fix planned	
1218	EXITINFO1[2] May Be Incorrectly Set When GMET (Guest Mode Execute Trap extension) is Enabled	No fix planned	
1225	MCA_STATUS_CS[ErrorCode] May Be Incorrect After Some Machine Check Errors	No fix planned	
1226	A Hardware Task Switch That Encounters Shadow Stack Errors May Cause a System Hang	No fix planned	
1227	A CPL3 to CPL3 Hardware Task Switch May Result in an Unexpected #CP	No fix planned	
1237	MCA_STATUS_LS[ExtErrorCode] May Contain Incorrect Code After a Store Queue Address Fatal Parity Error	No fix planned	
1238	MSRC001_1030[IbsIcMiss] May be Incorrect	No fix planned	
1244	EXITINFO1[2] May Be Incorrectly Set When Supervisor Shadow Stack Check is Enabled	No fix planned	
1257	MCA_STATUS_LS[Overflow] May Be Spuriously Set on DC_DATA_VICTIM or DC_DATA_LOAD Error	No fix planned	
1260	Processor May Log an Error with Incorrect Address in MCA_ADDR_LS	No fix planned	
1266	Retired MMX/FP Instruction Counter May Fail to Count MMX Stores	No fix planned	
1267	MSRC000_00E9[IRPerfCount] May Overcount	No fix planned	
1275	Software Using a Non-Canonicalized SSP (Shadow Stack Pointer) May Encounter a #GP Fault on a Different Instruction Than Expected	No fix planned	
1276	IOMMU May Fail to Abort Pre-Translated Request With Bits Erroneously Set in Reserved Field	No fix planned	
1277	IOMMU May Mishandle Fault on Skipped Page Directory Entry Levels	No fix planned	
1278	Some Features Are Not Available When SNP (Secure Nested Paging) Support Is Enabled	No fix planned	
1285	Processor May Fail to Report IBS (Instruction Based Sampling) Performance Sample	No fix planned	
1287	PMCx0AA[Source of Op Dispatched From Decoder] Events Will Not Be Counted	No fix planned	
1288	A VDPPS Instruction May Fail to Record a Masked Exception in the MXCSR Register	No fix planned	

Table 6. Cross-Reference of Processor Revision to Errata (continued)

No.	Errata Description	CUID Fn0000_0001_EAX	
		00A00F11h (Milan-B1)	00A00F12h (Milan-B2)
1292	Certain Performance Counters For Retire Based Events May Overcount	No fix planned	
1293	Data in IBS_OP_DATA2 and IBS_OP_DATA3 May be Inaccurate	No fix planned	
1294	xHCI Controller May Drop Data of an Isochronous TD (Transfer Descriptor) During Isochronous Transfer	No fix planned	
1295	Thread in Pending State From MONITOR or MONITORX May Become Unresponsive When Another Thread Invalidates the Monitored Address Using INVLPGB	No fix planned	
1296	Processor May Fail to Generate #GP on Incorrectly Programmed Host Save Address When Executing VMRUN	No fix planned	
1297	Processor May Fail To Take a #DB Exception on a Misaligned Store	No fix planned	
1298	Processor May Take Spurious #PF or Record Incorrect Error After DC_TAG_LOAD or DC_TAG_STORE Error	No fix planned	
1305	AHCI Controller Ignores COMINIT During HP6: HR_AwaitAlign State	No fix planned	
1307	Task Switches May Cause Unexpected #CP (Control Protection) Exception	No fix planned	
1308	Guests With AVIC (Advanced Virtual Interrupt Controller) Enabled May Not Be Able to Program Interrupt Controller	No fix planned	
1309	#VC (VMM Communication) Exception May Return Wrong Value in CR2	No fix planned	
1321	Processor May Generate Spurious #GP(0) Exception on WRMSR Instruction	No fix planned	
1322	PMCx0A9 May Undercount Op Queue Empty Events	No fix planned	
1329	Hypervisor With Encrypted SVM (Secure Virtual Machine) Related Pages May Experience Unpredictable Behavior If Guest Enables LBR (Last Branch Record)	X	
1330	SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) Guests With Last Branch Record Enabled May Experience Incorrect Alternate Injection Behavior	X	
1344	Processor May Fail to Take #VMEXIT(NPF) When SNP-Active (Secure Nested Paging) Guest Writes an Improperly Configured Page	No fix planned	
1345	RMPUPDATE or PSMASH Instructions May Speculatively Access the Page After End of RMP (Reverse Map Table)	No fix planned	
1346	RMPUPDATE May Incorrectly Return FAIL_OVERLAP	No fix planned	
1347	Instruction Cache L1TLB Page Size Reports Incorrectly	No fix planned	
1348	PMCx1CF May Undercount Operations Tagged by IBS (Instruction Based Sampling)	No fix planned	
1353	Write to Shared Core::X86::Msr::CpuWdtCfg From One Thread May Incorrectly Revert to Older Value	No fix planned	
1361	Processor May Hang When Switching Between Instruction Cache and Op Cache	X	
1363	Persistent Stream of APIC Register Accesses May Cause System to Hang or Reset	No fix planned	
1365	Fatal Error May Log Incorrect Location Information	No fix planned	
1373	RET or IRET to 32-Bit Mode When Upper 32 Bits of SSP (Shadow Stack Pointer) Are Not Equal to Zero May Cause Unexpected #GP	No fix planned	
1378	Secure TSC (Secure Time Stamp Counter) May Not Be Supported	No fix planned	
1379	VMSA (Virtual Machine Save Area) Register Protection May Not Be Supported	No fix planned	
1380	Processor May Speculatively Access Non-Speculative Memory or Cache Non-Cacheable Memory	No fix planned	
1381	Processor May Hang When Coherency Probe Hits Instruction Cache Line While Evicted	No fix planned	
1386	XSAVES Instruction May Fail to Save XMM Registers to the Provided State Save Area	No fix planned	
1390	PCIe® Port May Hang During Hot-plug Removal Event	No fix planned	
1391	Core::X86::Msr::IBS_OP_DATA3[IbsDcMissLat] May Overcount	No fix planned	

Table 6. Cross-Reference of Processor Revision to Errata (continued)

No.	Errata Description	CPUID Fn0000_0001_EAX	
		00A00F11h (Milan-B1)	00A00F12h (Milan-B2)
1394	Processor May Cause Unexpected Collisions on SMBUS (System Management Bus)	No fix planned	
1407	Processor May Signal Unexpected Fatal IF MCA Error	No fix planned	
1415	Processor Undergoing C-State Change May Signal Unexpected Fatal IF MCA Error	No fix planned	
1433	Processor May Log Unexpected LS MCE Error	No fix planned	
1450	Processor May Signal a Fatal LS MCA Error	No fix planned	
1455	PCID-Based INVLPGB May Fail to Flush Global Translations	No fix planned	
1464	32-Byte Misaligned Supervisor Shadow Stack Pointer or Supervisor Shadow Stack in Non-WB Memory May Result in a Non-Restartable Guest	No fix planned	
1467	Unexpected #PF for Hypervisor Write to 2M or 1G Page When SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) Enabled	No fix planned	
1476	Certain PCIe® Loads May Cause PCIe Completion Timeout	No fix planned	
1488	Translation Agent May Fail to Update Dirty Bit in Page Table Entry for Address Translation Request	No fix planned	
1489	IOMMU Unpinned Mode Is Not Supported When Secure Address Translation Service Is Enabled	X	

Cross-Reference of Errata to Package Type

This table cross-references the errata to each package type. "X" signifies that the erratum applies to the package type. An empty cell signifies that the erratum does not apply. An erratum may not apply to a package type due to a specific characteristic of the erratum, or it may be due to the affected silicon revision(s) not being used in this package.

Table 7. Cross-Reference of Errata to Package Type

Errata	Package
	SP3
1155	X
1159	X
1160	X
1169	X
1193	X
1197	X
1200	X
1211	X
1216	X
1218	X
1225	X
1226	X
1227	X
1237	X
1238	X
1244	X
1257	X
1260	X
1266	X
1267	X
1275	X
1276	X
1277	X
1278	X
1285	X
1287	X
1288	X
1292	X
1293	X
1294	X

Table 7. Cross-Reference of Errata to Package Type
(continued)

Errata	Package
	SP3
1295	X
1296	X
1297	X
1298	X
1305	X
1307	X
1308	X
1309	X
1321	X
1322	X
1329	X
1330	X
1344	X
1345	X
1346	X
1347	X
1348	X
1353	X
1361	X
1363	X
1365	X
1373	X
1378	X
1379	X
1380	X
1381	X
1386	X
1390	X
1391	X
1394	X
1407	X
1415	X
1433	X
1450	X
1455	X

**Table 7. Cross-
Reference of Errata to
Package Type
(continued)**

Errata	Package
	SP3
1464	X
1467	X
1476	X
1488	X
1489	X

Cross-Reference of Errata to Processor Segments

This table cross-references the errata to each processor segment. "X" signifies that the erratum applies to the processor segment. An empty cell signifies that the erratum does not apply. An erratum may not apply to a processor segment due to a specific characteristic of the erratum, or it may be due to the affected silicon revision(s) not being used in this processor segment.

Table 8. Cross-Reference of Errata to Processor Segments

Errata	Processor Segment
	EPYC 7003 Generation Processors
1155	X
1159	X
1160	X
1169	X
1193	X
1197	X
1200	X
1211	X
1216	X
1218	X
1225	X
1226	X
1227	X
1237	X
1238	X
1244	X
1257	X
1260	X
1266	X
1267	X
1275	X
1276	X
1277	X
1278	X
1285	X
1287	X
1288	X
1292	X

Table 8. Cross-Reference of Errata to Processor Segments (continued)

Errata	Processor Segment
	EPYC 7003 Generation Processors
1293	X
1294	X
1295	X
1296	X
1297	X
1298	X
1305	X
1307	X
1308	X
1309	X
1321	X
1322	X
1329	X
1330	X
1344	X
1345	X
1346	X
1347	X
1348	X
1353	X
1361	X
1363	X
1365	X
1373	X
1378	X
1379	X
1380	X
1381	X
1386	X
1390	X
1391	X
1394	X
1407	X
1415	X

**Table 8. Cross-
Reference of Errata to
Processor Segments
(continued)**

Errata	Processor Segment
	EPYC 7003 Generation Processors
1433	X
1450	X
1455	X
1464	X
1467	X
1476	X
1488	X
1489	X

1155 DMA or Peer-to-peer Accesses Using Guest Physical Addresses (GPAs) May Cause IOMMU Target Abort

Description

In systems where:

- Virtualization is enabled, and
- IOMMU is in pass-through mode

DMA or peer-to-peer accesses using Guest Physical Addresses (GPAs) occurring within the regions defined below trigger a target abort.

- 0x00FD_0000_0000->0x00FD_F8FF_FFFF, or
- 0x00FD_F910_0000->0x00FD_F91F_FFFF, or
- 0x00FD_FB00_0000->0x00FD_FFFF_FFFF

Potential Effect on System

A DMA device will receive a target abort from the IOMMU.

Suggested Workaround

System software must mark the following block of memory as reserved:

- FD_0000_0000 -> FD_FFFF_FFFF

Fix Planned

No fix planned

1159 Writes to Base Frequency Register May be Ignored

Description

If the base frequency register, (MSRC001_0064[CpuDfsId], MSRC001_0064[CpuFid]):

- is programmed to a lower frequency than the default reset value, and
- the default base frequency is not a multiple of 100 MHz

then subsequent writes to the register that are greater than the next lower multiple of 100 MHz may be ignored.

Potential Effect on System

Software may report an incorrect value of base frequency.

Suggested Workaround

None

Fix Planned

No fix planned

1160 SdpParity and XiVictimQueue Mask Bits Incorrectly Mask Additional Errors

Description

If MCA::L3::MCA_CTL_MASK_L3[5] (SdpParity) is set then errors logged in MCA_STATUS_L3 that set MCA_STATUS_L3[ErrorCodeExt]=0x5 are correctly masked, and some system read data errors logged in MCA_STATUS_LS or MCA_STATUS_IF are masked incorrectly.

If MCA::L3::MCA_CTL_MASK_L3[6] (XiVictimQueue) is set then errors logged in MCA_STATUS_L3 that set MCA_STATUS_L3[ErrorCodeExt]=0x6 are correctly masked, and some system read data errors logged in MCA_STATUS_LS are masked incorrectly.

Potential Effect on System

Some system read data errors logged in MCA_STATUS_LS or MCA_STATUS_IF may fail to be detected.

Suggested Workaround

Do not program MCA::L3::MCA_CTL_MASK_L3[5] or MCA::L3::MCA_CTL_MASK_L3[6] to 1b.

Fix Planned

No fix planned

1169 PCIe® Error Masking May Fail to Mask Errors

Description

If MCA_CTL_MASK_NBIO[PCIE_sideband] is programmed to 0b, then PCIe® error masking, including Uncorrectable Error Mask and Correctable Error Mask, will not mask errors.

Potential Effect on System

Masked errors will incorrectly be reported to the system.

Suggested Workaround

Program MCA_CTL_MASK_NBIO[PCIE_sideband] to 1b.

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1193 Page Remapping Without Invalidation May Cause Missed Detection of Self-Modifying Code

Description

Under a highly specific and detailed set of internal timing conditions, if a page table entry that has the Accessed bit set has its Page Frame Number (Physical-Page Base Address) updated without first updating the translation to have a permission violation for instruction fetch and then invalidating the table entry (INVLPG or INVLPGB +DVMSYNC), the processor may:

- Execute stale instructions in the presence of self-modifying code or cross-modifying code for an instruction using the old translation.

Potential Effect on System

Processor may execute stale instructions in the presence of self-modifying code or cross-modifying code.

Suggested Workaround

When remapping a page that is used for code fetch, software should update the translation to have a permission violation and invalidate the page table entry prior to remapping.

Fix Planned

No fix planned

1197 IBS (Instruction Based Sampling) Register State May be Incorrect After Restore From CC6

Description

If an IBS (Instruction Based Sampling) interrupt is asserted, but the processor begins entry into the CC6 state before the IBS interrupt is processed, then the IBS register state may be incorrect after the restore from CC6.

Potential Effect on System

Debugger or performance monitoring software that uses IBS functionality may encounter inaccurate data.

Suggested Workaround

Software should ignore IBS samples if `MSRC001_1031 == 0000_0000_0000_0000h`.

Fix Planned

No fix planned

1200 xHCI Host May Hang If Full Speed or High Speed USB Hub is Connected

Description

xHCI Host controller may hang if:

- A high speed or full speed flash device is connected to the host, and
- A high speed or full speed hub is connected to the host, and
- An active device is connected to the hub.

Potential Effect on System

xHCI Host controller may hang.

Suggested Workaround

System software may contain the workaround for this erratum.

Program USB0x00C60C[12:9] to 0001b.

Program USB1x00C60C[12:9] to 0001b.

Program USB0x00C608[6:5] to 00b.

Program USB1x00C608[6:5] to 00b.

Fix Planned

No fix planned

1211 PCIe Incorrectly Logs TLP Prefix AER (Advanced Error Reporting) Register After Egress Blocking Error

Description

When a PCIe TLP prefix egress blocking error occurs in an ATS (Address Translation Service) invalidation message or PRG (Page Request Group) response message, a TPH (TLP Processing Hint) prefix will be incorrectly be logged in the TLP Prefix AER (Advanced Error Reporting) register instead of PASID (Process Address Space ID).

Potential Effect on System

Software reading the TLP Prefix AER register may encounter incorrect information.

Suggested Workaround

Ignore Bit 24 of the TLP Prefix AER after a PCIe egress blocking error in an ATS invalidation message or PRG response message.

Fix Planned

No fix planned

1216 IOMMU May Not Re-Walk Page Tables on a Present/Permission Fault

Description

Under a highly specific and detailed set of internal timing conditions, the IOMMU may not re-walk page tables on a present/permission fault.

Potential Effect on System

An IO device may see an unexpected completer abort.

Suggested Workaround

System software may contain the workaround for this erratum.

Program IOMMUL2B0x00000150[16] to 1b, and

Program IOMMUL2B1x00000150[16] to 1b, and

Program IOMMUL2B2x00000150[16] to 1b, and

Program IOMMUL2B3x00000150[16] to 1b.

Fix Planned

No fix planned

1218 EXITINFO1[2] May Be Incorrectly Set When GMET (Guest Mode Execute Trap extension) is Enabled

Description

EXITINFO1[2] (User/Supervisor bit) may incorrectly be one during a nested page fault if GMET (Guest Mode Execute Trap extension) is enabled.

Potential Effect on System

Software may not be able to determine whether a fault was a GMET fault or an NX fault based on EXITINFO1.

Suggested Workaround

Software must read the relevant VMCB to determine whether a fault was a GMET fault or an NX fault.

Fix Planned

No fix planned

1225 MCA_STATUS_CS[ErrorCode] May Be Incorrect After Some Machine Check Errors

Description

For some machine check errors, the MCA_STATUS_CS[ErrorCode] memory transaction type (RRRR) field incorrectly contains the value 0101b (Instruction Fetch) for all transaction types.

This erratum affects errors logged with the following MCA_STATUS_CS[ErrorCodeExt]:

- FTI_ILL_REQ
- FTI_ADDR_VIOL
- FTI_SEC_VIOL
- FTI_ILL_RSP
- FTI_RSP_NO_MTCH
- SPF_PRT_ERR
- SDP_RSP_NO_MTCH
- SDP_UNEXP_RETRY
- CNTR_OVFL
- CNTR_UNFL

Potential Effect on System

None expected. Software is expected to primarily rely on MCA_STATUS_CS[ErrorCodeExt] to identify errors.

Suggested Workaround

None. Software should use MCA_STATUS_CS[ErrorCodeExt] to identify errors and ignore MCA_STATUS_CS[ErrorCode] RRRR value for the error types listed above.

Fix Planned

No fix planned

1226 A Hardware Task Switch That Encounters Shadow Stack Errors May Cause a System Hang

Description

A hardware task switch that encounters a shadow stack error condition may result in a #GP instead of a #TS. A hardware task switch that encounters a fault during shadow stack accesses may result in TR (Task Register) being incorrectly loaded.

Potential Effect on System

System may hang or reset.

Suggested Workaround

None

Fix Planned

No fix planned

1227 A CPL3 to CPL3 Hardware Task Switch May Result in an Unexpected #CP

Description

A CPL3 to CPL3 hardware task switch may result in the SSP (Shadow Stack Pointer) being corrupted.

Potential Effect on System

System may encounter an unexpected #CP.

Suggested Workaround

None

Fix Planned

No fix planned

1237 MCA_STATUS_LS[ExtErrorCode] May Contain Incorrect Code After a Store Queue Address Fatal Parity Error

Description

Under a highly specific and detailed set of internal timing conditions, an error that should be logged with MCA_STATUS_LS[ExtErrorCode]=11 (store queue address fatal parity error) may be incorrectly logged with MCA_STATUS_LS[ExtErrorCode]=23 (other store data fatal error).

Potential Effect on System

Diagnostic software may encounter incorrect information in MCA_STATUS_LS[ExtErrorCode] after a fatal error.

Suggested Workaround

None

Fix Planned

No fix planned

1238 MSRC001_1030[IbsIcMiss] May be Incorrect

Description

The processor may set MSRC001_1030[IbsIcMiss] as an instruction cache miss when the fetch was an instruction cache hit.

Potential Effect on System

MSRC001_1030[IbsIcMiss] may be incorrect.

Suggested Workaround

None.

Ignore MSRC001_1030[IbsIcMiss].

Fix Planned

No fix planned

1244 EXITINFO1[2] May Be Incorrectly Set When Supervisor Shadow Stack Check is Enabled

Description

EXITINFO1[2] (User/Supervisor bit) may incorrectly be 1b during a nested page fault if Supervisor Shadow Stack Check (Virtual Machine Control Block offset 90h bit 4) is enabled.

Potential Effect on System

Software may not be able to determine whether a fault was a Supervisor Shadow Stack Check or not based on EXITINFO1[2].

Suggested Workaround

Software may examine the following to determine the type of the fault:

- EXITINFO1[37] and EXITINFO1[6]
- CPL (Current Privilege Level) from VMCB (Virtual Machine Control Block)
- Guest instruction bytes from VMCB (WRUSS instruction is a user access regardless of CPL)
- Nested page table

Fix Planned

No fix planned

1257 MCA_STATUS_LS[Overflow] May Be Spuriously Set on DC_DATA_VICTIM or DC_DATA_LOAD Error

Description

Under a highly specific and detailed set of internal timing conditions, MCA_STATUS_LS[Overflow] may be spuriously set when a DC_DATA_VICTIM or DC_DATA_LOAD error is logged.

Potential Effect on System

None

Suggested Workaround

None

Fix Planned

No fix planned

1260 Processor May Log an Error with Incorrect Address in MCA_ADDR_LS

Description

Under a highly specific and detailed set of internal timing conditions, when logging an MCA error due to poison consumption (MCA_STATUS_LS[ExtErrorCode]=DC_DATA_LOAD and MCA_STATUS_LS[Poison]=1), the processor may log an incorrect address in MCA_ADDR_LS.

Potential Effect on System

System software may incorrectly identify which address contains poisoned data.

Suggested Workaround

System software should use the deferred error information associated with poison creation to identify the address containing poisoned data.

Fix Planned

No fix planned

1266 Retired MMX/FP Instruction Counter May Fail to Count MMX Stores

Description

If PMCx0CB[MmxInstr] is programmed to 1b, the processor may experience sampling inaccuracies that cause the PMCx0CB [Retired MMX/FP Instructions] performance counter to fail to count MMX stores.

Potential Effect on System

Performance monitoring software may undercount retired MMX/FP instructions.

Suggested Workaround

None

Fix Planned

No fix planned

1267 MSRC000_00E9[IRPerfCount] May Overcount

Description

The processor may experience sampling inaccuracies that cause MSRC000_00E9[IRPerfCount] to overcount certain instruction fusion cases.

Potential Effect on System

Inaccuracies in performance monitoring software may be experienced.

Suggested Workaround

If absolute accuracy is required, then use PMCx0C0 instead.

Fix Planned

No fix planned

1275 Software Using a Non-Canonicalized SSP (Shadow Stack Pointer) May Encounter a #GP Fault on a Different Instruction Than Expected

Description

The processor may not canonicalize the SSP before saving for far transfers when:

- EFER.LMA=1, or
- The far transfer is a SYSCALL

The processor may not fault during a far transfer if the target SSP is non-canonical.

Potential Effect on System

None expected. Software using a non-canonicalized SSP (Shadow Stack Pointer) may encounter a #GP fault on a different instruction than expected.

Suggested Workaround

The operating system or hypervisor may contain the workaround for this erratum.

Fix Planned

No fix planned

1276 IOMMU May Fail to Abort Pre-Translated Request With Bits Erroneously Set in Reserved Field

Description

The IOMMU fails to abort pre-translated request with bits erroneously set in the reserved host PDE[60:52] (Page Directory Entry) or reserved host PTE[58:52] (Page Table Entry) fields when sATS (Secure Address Translation Service) is enabled for the device making the request.

Potential Effect on System

None expected. Properly coded software will avoid programming reserved bits in the host PDE or host PTE structures.

Suggested Workaround

Do not program bits in:

- the reserved host PDE[60:52] (Page Directory Entry) field, or
- the reserved host PTE[58:52] (Page Table Entry) field

Fix Planned

No fix planned

1277 IOMMU May Mishandle Fault on Skipped Page Directory Entry Levels

Description

When Guest Page Tables and Nested Page Tables are enabled, if a nested page table walk skips a PDE (Page Directory Entry) level when the virtual address bits are non-zero, the IOMMU may fail to abort the request, and fail to generate an IO page fault.

Potential Effect on System

None expected. Properly coded software will program the virtual address bits associated with a skipped page level to all zero.

Suggested Workaround

Program the virtual address bits associated with a skipped page level to all zero.

Fix Planned

No fix planned

1278 Some Features Are Not Available When SNP (Secure Nested Paging) Support Is Enabled

Description

The following features are not available when SNP support is enabled:

- IOMMU Guest Paging (GTSup) Two-level Guest Address Translation
- IOMMU SATS (SATSSup) Secure Address Translation Service
- IOMMU AVIC (GAMSup) Guest Virtual APIC Interrupt Controller
-
-

Potential Effect on System

The features listed above are not available when SNP support is enabled.

Suggested Workaround

None

Fix Planned

No fix planned

1285 Processor May Fail to Report IBS (Instruction Based Sampling) Performance Sample

Description

If Core::X86::Msr::IBS_OP_CTL[IbsOpEn] is programmed to 1b, the processor may:

- fail to signal an interrupt after capturing an IBS sample and
- erroneously leave Core::X86::Msr::IBS_OP_CTL[IbsOpVal] equal to 1b, and
- fail to capture further samples until Core::X86::Msr::IBS_OP_CTL[IbsOpEn] is programmed to 1b again.

Potential Effect on System

Profiling tools using IBS may fail to capture some IBS records.

Suggested Workaround

If IBS and PMC (Performance Monitor Counter) overflow interrupts are both configured to be delivered as NMIs (Non-Maskable Interrupts), then software may program a PMC to create periodic interrupts so the IBS handler can detect the existence of the sample.

Example 1: If using Linux perf tool with IBS configured to use cycle-based sampling:perf record -a -e ibs_op/cnt_ctl=0/cycles <workload>

Example 2: If using Linux perf tool with IBS configured to use micro-op-based sampling:perf record -a -e ibs_op/cnt_ctl=1/cpu/event=0xc1/ <workload>

An alternate workaround is to set the BIOS Setup Option "IBS hardware workaround." This option is for profiling code that uses IBS optimizations and is not recommended for production systems because system performance may be negatively impacted.

Fix Planned

No fix planned

1287 PMCx0AA[Source of Op Dispatched From Decoder] Events Will Not Be Counted

Description

If a core performance monitor counter is programmed to select PMCx0AA[Source of Op Dispatched From Decoder], the counter will not count the selected events.

Potential Effect on System

Performance monitoring software may undercount "Source of Op Dispatched From Decoder" events.

Suggested Workaround

None

Fix Planned

No fix planned

1288 A VDPPS Instruction May Fail to Record a Masked Exception in the MXCSR Register

Description

A 256-bit VDPPS instruction will fail to record a masked exception in the MXCSR register when:

- An unmasked exception is detected on one 128-bit section in the addition phase of the instruction, and
- A masked exception is detected on the other 128-bit section in the addition phase of the instruction.

Potential Effect on System

None expected.

Suggested Workaround

None

Fix Planned

No fix planned

1292 Certain Performance Counters For Retire Based Events May Overcount

Description

The processor may experience sampling inaccuracies that cause the following performance counters to overcount retire-based events.

- PMCx0C0 [Retired Instructions]
- PMCx0C1 [Retired Uops]
- PMCx0C2 [Retired Branch Instructions]
- PMCx0C3 [Retired Branch Instructions Mispredicted]
- PMCx0C4 [Retired Taken Branch Instructions]
- PMCx0C5 [Retired Taken Branch Instructions Mispredicted]
- PMCx0C8 [Retired Near Returns]
- PMCx0C9 [Retired Near Returns Mispredicted]
- PMCx0CA [Retired Indirect Branch Instructions Mispredicted]
- PMCx0CC [Retired Indirect Branch Instructions]
- PMCx0D1 [Retired Conditional Branch Instructions]
- PMCx1C7 [Retired Mispredicted Branch Instructions due to Direction Mismatch]
- PMCx1D0 [Retired Fused Branch Instructions]

Potential Effect on System

Inaccuracies in performance monitoring software may be experienced.

Suggested Workaround

To count the non-FP affected PMC events correctly:

- Use Core::X86::Msr::PERF_CTL2 to count the events, and
- Program Core::X86::Msr::PERF_CTL2[43] to 1b, and
- Program Core::X86::Msr::PERF_CTL2[20] to 0b.

An alternate workaround that enables the capture of multiple affected PMC events simultaneously is to set the BIOS Setup Option "IBS hardware workaround." This option is not recommended for production systems because system performance may be negatively impacted.

Fix Planned

No fix planned

1293 Data in IBS_OP_DATA2 and IBS_OP_DATA3 May be Inaccurate

Description

If either of the following conditions are met:

- Core::X86::Msr::IBS_OP_DATA3[DcMissNoMabAlloc] = 1b, or
- Core::X86::Msr::IBS_OP_DATA3[IbsSwPf] = 1b

Then the data in the following registers may be inaccurate:

- Core::X86::Msr::IBS_OP_DATA2
- Core::X86::Msr::IBS_OP_DATA3[IbsOpDcMissOpenMemReqs]
- Core::X86::Msr::IBS_OP_DATA3[IbsL2Miss]

Potential Effect on System

Performance monitoring software may encounter incorrect data in IBS_OP_DATA2 or IBS_OP_DATA3.

Suggested Workaround

None.

Performance monitoring software should ignore the affected registers in samples with IBS_OP_DATA3[DcMissNoMabAlloc] = 1b or IBS_OP_DATA3[IbsSwPf] = 1b.

Fix Planned

No fix planned

1294 xHCI Controller May Drop Data of an Isochronous TD (Transfer Descriptor) During Isochronous Transfer

Description

When an Evaluate Context Command modifies the Max Exit Latency value when an Isochronous transfer is in progress, the xHCI controller may drop the data of an Isochronous TD of the endpoint associated with the Device Slot targeted by the Evaluate Context Command. This may result in the xHCI issuing an MSE (Missed Service Error).

Potential Effect on System

Isochronous Audio or Video transfers may experience momentary data loss within a 750 microsecond timeout window, after which isochronous transfer will resume.

Suggested Workaround

None

Fix Planned

No fix planned

1295 Thread in Pending State From MONITOR or MONITORX May Become Unresponsive When Another Thread Invalidates the Monitored Address Using INVLPGB

Description

The processor may fail to exit the monitor event pending state entered by a MONITOR or MONITORX instruction under the following conditions:

- An MWAIT or MWAITX instruction is executed while the processor is in the event pending state.
- A different processor or processor thread executes an INVLPGB instruction that invalidates the translation for the address range established by the MONITOR or MONITORX instruction, and
- The processor that executed the INVLPGB then executes a TLBSYNC instruction, before the first processor exits the monitor event pending state.

Under the conditions specified above, the first processor may remain in the event pending state until an unmasked interrupt is seen or a write happens to the physical address that belonged to the linear address of the MONITOR/MONITORX instruction when executed.

An MWAITX instruction with a timer interval programmed would wake up at the expiration of the timer interval.

Potential Effect on System

A thread may become unresponsive until the next interrupt.

Suggested Workaround

None

Fix Planned

No fix planned

1296 Processor May Fail to Generate #GP on Incorrectly Programmed Host Save Address When Executing VMRUN

Description

The processor may fail to generate a #GP fault when:

- A VMRUN instruction is executed, and
- Core::X86::Msr::VM_HSAVE_PA = 8_0000_0000_0000h, and
- Core::X86::Msr::SYS_CFG[SecureNestedPagingEn] = 1b

Potential Effect on System

None

Suggested Workaround

None

Fix Planned

No fix planned

1297 Processor May Fail To Take a #DB Exception on a Misaligned Store

Description

Under a highly specific and detailed set of internal timing conditions, the processor may fail to take a #DB exception when a store that is misaligned on a 4K address boundary matches a data breakpoint on the portion of the store that is after the 4K boundary crossing.

Potential Effect on System

Missed data breakpoint.

Suggested Workaround

None

Fix Planned

No fix planned

1298 Processor May Take Spurious #PF or Record Incorrect Error After DC_TAG_LOAD or DC_TAG_STORE Error

Description

Under a highly detailed and specific set of internal timing conditions, after the processor detects a DC_TAG_LOAD or DC_TAG_STORE error, one of the following may occur:

- The processor may take a spurious #PF (Page Fault Exception) for an address that is not faulting prior to handling the MCA error.
- The processor may record the wrong information in MCA_ADDR_LS and MCA_STATUS_LS[ErrorCodeExt].

Potential Effect on System

Incorrect error reporting and/or spurious #PF.

Suggested Workaround

None

Fix Planned

No fix planned

1305 AHCI Controller Ignores COMINIT During HP6: HR_AwaitAlign State

Description

In HP6: HR_AwaitAlign state, while the AHCI controller is awaiting valid ALIGN patterns from connected SATA device, it will not respond to COMINIT issued by the connected SATA device.

Potential Effect on System

If the attached SATA device sends COMINIT instead of valid ALIGN patterns in HP6:HR_AwaitAlign state, the AHCI controller will time out awaiting valid ALIGN patterns. Consequently the AHCI controller will re-initiate Out-of-band signaling sequence at the next highest supported speed. This may result in the attached SATA device running at the lower speed.

Suggested Workaround

None

Fix Planned

No fix planned

1307 Task Switches May Cause Unexpected #CP (Control Protection) Exception

Description

Under a highly detailed sequence of internal timing conditions, the processor may fail to push CS (Code Segment), RIP (Linear Instruction Pointer), and SSP (Shadow Stack Pointer) onto the shadow stack when executing a CALL FAR instruction if the following conditions are also met:

- Long mode is disabled (EFER.LMA = 0).
- The CALL FAR instruction is executed while processor is running at CPL0.
- The segment-descriptor type referenced by the CALL FAR instruction is an available TSS (task state segment).
- User shadow stack is disabled and supervisor shadow stack is enabled (U_CET[SH_STK_EN] = 0 and S_CET[SH_STK_EN] = 1).

Potential Effect on System

None expected.

If task switches are used, unexpected shadow stack mismatch leading to a #CP (control protection) exception may occur.

Suggested Workaround

None.

Operating systems using the shadow stack features should avoid using task switches.

Fix Planned

No fix planned

1308 Guests With AVIC (Advanced Virtual Interrupt Controller) Enabled May Not Be Able to Program Interrupt Controller

Description

When AVIC (Advanced Virtual Interrupt Controller) is enabled, the processor may fail to redirect accesses to the AVIC backing page if the system PA (Physical Address) for APIC_BAR (Advanced Programmable Interrupt Controller Base Address Register) in the nested page table is an MMIO (Memory Mapped IO) address.

Potential Effect on System

Guests with AVIC enabled may not be able to program the interrupt controller.

Suggested Workaround

Ensure that the system PA for APIC_BAR in the nested page table is not an MMIO address.

Fix Planned

No fix planned

1309 #VC (VMM Communication) Exception May Return Wrong Value in CR2

Description

Under a highly specific and detailed set of internal timing conditions, a #VC (VMM Communication) exception may return the wrong value in CR2 under the following conditions:

- The exception is taken on an RMPADJUST instruction in a Secure Nested Paging (SNP)-active guest.
- The error code is PAGE_NOT_VALIDATED.
- The RMP (Reverse Map Table) entry accessed by the RMPADJUST instruction is concurrently accessed by a PVALIDATE instruction from the same guest running on a different processor thread.

Potential Effect on System

The guest may incorrectly indicate an error and/or terminate.

Suggested Workaround

None

Fix Planned

No fix planned

1321 Processor May Generate Spurious #GP(0) Exception on WRMSR Instruction

Description

The processor will generate a spurious #GP(0) exception on a WRMSR instruction if the following conditions are all met:

- The target of the WRMSR is the SYSCFG register.
- The write changes the value of Secure Nested Paging enable (SYSCFG.SNPEn) from 0 to 1.
- One of the threads that share the physical core has a non-zero value in the VM_HSAVE_PA MSR.

Potential Effect on System

Unexpected #GP(0) exception during processor boot.

Suggested Workaround

When enabling Secure Nested Paging, program VM_HSAVE_PA to 0h on both threads that share a physical core before setting SYSCFG.SNPEn.

Fix Planned

No fix planned

1322 PMCx0A9 May Undercount Op Queue Empty Events

Description

If a core performance monitor counter is programmed to select PMCx0A9[Op Queue Empty], the processor may experience sampling inaccuracies that cause some "Op Queue Empty" events not to be counted.

Potential Effect on System

Performance monitoring software may undercount "Op Queue Empty" events.

Suggested Workaround

None

Fix Planned

No fix planned

1329 Hypervisor With Encrypted SVM (Secure Virtual Machine) Related Pages May Experience Unpredictable Behavior If Guest Enables LBR (Last Branch Record)

Description

Under the following conditions, the processor may behave incorrectly:

- SME (Secure Memory Encryption) is enabled, and
- Hypervisor owned pages used for SVM (Secure Virtual Machine) functionality are encrypted, and
- A VMRUN instruction has been executed to enter an SEV-ES (Secure Encrypted Virtualization - Encrypted State) or SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) guest, and
- DbgCtl[LBR] (Last Branch Record) was 0 when the hypervisor started the guest with a VMRUN instruction, and
- The guest value of DbgCtl[LBR] was 1 when loaded from the VMSA (Virtual Machine Save Area).

The incorrect behavior may be one or more of the following:

- The processor may unexpectedly cause a VMEXIT or may fail to cause a VMEXIT.
- VMCB.GUEST_INSTR_BYTES may have an incorrect value on VMEXIT.
- If the nested page tables are in encrypted memory, nested page table walks may result in incorrect translations.
- If AVIC (Advanced Virtual Interrupt Controller) is enabled and the AVIC backing page is encrypted, AVIC functionality may be incorrect.

Potential Effect on System

Unpredictable system behavior

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1330 SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) Guests With Last Branch Record Enabled May Experience Incorrect Alternate Injection Behavior

Description

Under the following conditions, a SEV-SNP (Secure Encrypted Virtualization Secure Nested Paging) guest with Alternate Injection enabled may incorrectly behave as if Alternate Injection is disabled:

- DbgCtl[LBR] (Last Branch Record) was 0 when the hypervisor started the guest with a VMRUN instruction, and
- The guest value of DbgCtl[LBR] was 1 when loaded from the VMSA (Virtual Machine Save Area).

Potential Effect on System

Incorrect interrupt behavior for guests.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1344 Processor May Fail to Take #VMEXIT(NPF) When SNP-Active (Secure Nested Paging) Guest Writes an Improperly Configured Page

Description

The processor may fail to take a #VMEXIT(NPF) under the following conditions:

- An SNP-Active guest executes an instruction that performs a memory write or a masked write (VMASKMOV), and
- The page that is accessed by the memory operation is illegally configured to have VMPL (Virtual Machine Privilege Level) write permissions but not VMPL read permissions.

Potential Effect on System

None expected. Properly coded software will always enable VMPL read permissions if VMPL write permissions are enabled.

Suggested Workaround

None

Fix Planned

No fix planned

1345 RMPUPDATE or PSMASH Instructions May Speculatively Access the Page After End of RMP (Reverse Map Table)

Description

When executing the RMPUPDATE or PSMASH instructions, the processor may speculatively read from addresses that are in the next 4K page after the end of the RMP even if the memory type of that page does not allow speculative accesses.

Potential Effect on System

System may hang or reset.

Suggested Workaround

Always map a 4K page that allows speculative accesses after the end of the RMP.

Fix Planned

No fix planned

1346 RMPUPDATE May Incorrectly Return FAIL_OVERLAP

Description

RMPUPDATE may fail with FAIL_OVERLAP return code if multiple processors simultaneously execute RMPUPDATE to change the assignment of different pages within a 2MB region.

Potential Effect on System

A hypervisor may receive an unexpected FAIL_OVERLAP return code.

Suggested Workaround

Retry the operation.

Fix Planned

No fix planned

1347 Instruction Cache L1TLB Page Size Reports Incorrectly

Description

Core::X86::Msr::IBS_FETCH_CTL[IbsL1TlbPgSz] instruction cache L1TLB page size reports incorrectly. The processor instead reports as follows:

- IbsL1TlbPgSz = 00b: 4KB
- IbsL1TlbPgSz = 01b: 16KB
- IbsL1TlbPgSz = 10b: 2MB
- IbsL1TlbPgSz = 11b: 1GB

Potential Effect on System

Performance monitoring software may produce unexpected results.

Suggested Workaround

Software reading IbsL1TlbPgSz should refer to the table in the description.

Fix Planned

No fix planned

1348 PMCx1CF May Undercount Operations Tagged by IBS (Instruction Based Sampling)

Description

If PMCx1CF[IbsTaggedOps] is programmed to 1b, the PMCx1CF [Tagged IBS Ops] performance counter may undercount when sampling thread 1.

Potential Effect on System

Inaccuracies in performance monitoring software may be experienced.

Suggested Workaround

None

Fix Planned

No fix planned

1353 Write to Shared Core::X86::Msr::CpuWdtCfg From One Thread May Incorrectly Revert to Older Value

Description

A write to the shared Core::X86::Msr::CpuWdtCfg from one thread may incorrectly be reverted to an older value under the following conditions:

- The CpuWdtCfg write happens while the other thread is in HLT, and
- The CpuWdtCfg write happens between a CC6 state exit and a CC6 state entry, and
- The other thread remained in HLT for the entire time, and
- The next time the other thread wakes up from HLT, it wakes up after the writing thread woke up subsequent to a CC6 state exit.

Potential Effect on System

Incorrect watchdog timer functionality

Suggested Workaround

Program the CpuWdtCfg on one thread while the other thread is not in HLT.

Fix Planned

No fix planned

1361 Processor May Hang When Switching Between Instruction Cache and Op Cache

Description

Under a highly specific and detailed set of internal timing conditions, running a program with a code footprint that exceeds 32 KB may cause the processor to hang while switching between code regions that consistently miss the instruction cache and code regions that are contained within the Op Cache.

Potential Effect on System

System may hang or reset.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1363 Persistent Stream of APIC Register Accesses May Cause System to Hang or Reset

Description

A persistent stream of APIC register accesses, such as a write to the Task Priority Register (TPR), from one set of processor cores may prevent the system from completing a Message Signaled Interrupt (MSI) or APIC store or x2APIC register write operation from another I/O device or processor core that asserts an interrupt.

Potential Effect on System

System may hang or reset.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1365 Fatal Error May Log Incorrect Location Information

Description

Under a highly specific and detailed set of internal timing conditions, a fatal error logged with MCA::LS::MCA_STATUS_LS[ErrorCodeExt]=0xC:

- May log the incorrect location in MCA::LS::MCA_SYND_LS[ErrorInformation], or
- May be recorded for the other logical thread sharing the same core.

Potential Effect on System

The system may log errors erroneously.

Suggested Workaround

None

Fix Planned

No fix planned

1373 RET or IRET to 32-Bit Mode When Upper 32 Bits of SSP (Shadow Stack Pointer) Are Not Equal to Zero May Cause Unexpected #GP

Description

The processor will generate an unexpected #GP under the following conditions:

- There is a RET or IRET to a 32-bit mode when target CPL is 0, 1 or 2.
- Supervisor Shadow Stack is disabled.
- SSP (Shadow Stack Pointer) [63:32] is not equal to zero.
- User Shadow Stack is enabled.

Potential Effect on System

Unexpected #GP

Suggested Workaround

None

Fix Planned

No fix planned

1378 Secure TSC (Secure Time Stamp Counter) May Not Be Supported

Description

Support for SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) may not include the Secure TSC feature. If Secure TSC is not supported, then reading `CPUID_Fn8000001F_EAX[SecureTsc]` will (correctly) return 0b to indicate this lack of support.

Potential Effect on System

The Secure TSC feature may not be supported.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1379 VMSA (Virtual Machine Save Area) Register Protection May Not Be Supported

Description

Support for SEV-ES (Secure Encrypted Virtualization - Encrypted State) may not include the VMSA (Virtual Machine Save Area) register protection feature. If VMSA is not supported, then reading `CPUID_Fn8000001F_EAX[VMSARegProt]` will (correctly) return 0b to indicate this lack of support.

Potential Effect on System

The VMSA register protection feature may not be supported.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1380 Processor May Speculatively Access Non-Speculative Memory or Cache Non-Cacheable Memory

Description

Under a highly specific and detailed set of internal timing conditions, the processor may do one of the following:

- Prefetch data into the cache from a 4KB page that has a non-cacheable memory type or for which the process executing on the processor does not have permissions.
- Speculatively make prefetch accesses to memory addresses from a 4KB page that does not allow speculative memory accesses or for which the process executing on the processor does not have permissions.

In the case where the processor does not have permissions, the executing process does not get access to the stored values that were prefetched into the cache.

The incorrect behavior can only happen for 4KB pages that are adjacent in the physical address space to a 4KB page with WB memory type to which the process executing on the processor does have permissions. The adjacent 4KB WB memory page must also be within the same 2MB aligned region as the incorrect access.

Potential Effect on System

None expected. Data may be cached for addresses not marked as cacheable in the page tables.

Suggested Workaround

None

Fix Planned

No fix planned

1381 Processor May Hang When Coherency Probe Hits Instruction Cache Line While Evicted

Description

Under a highly specific and detailed set of internal timing conditions, a processor may hang while filling a line into the instruction cache if a coherency probe is simultaneously received and the address of the probe is the same as the address of the line that is being evicted from the instruction cache.

Potential Effect on System

System may hang or reset.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1386 XSAVES Instruction May Fail to Save XMM Registers to the Provided State Save Area

Description

The XSAVES instruction may fail to save XMM registers to the provided state save area if all of the following are true:

- All XMM registers were restored to the initialization value by the most recent XRSTORS instruction because the XSTATE_BV[SSE] bit was clear.
- The state save area for the XMM registers does not contain the initialization state.
- The value in the XMM registers match the initialization value when the XSAVES instruction is executed.
- The MXCSR register has been modified to a value different from the initialization value since the most recent XRSTORS instruction.

Potential Effect on System

A process may observe a non-zero value in an XMM register that should be zero.

This observed value will be a previous XMM value from the same process.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1390 PCIe® Port May Hang During Hot-plug Removal Event

Description

A PCIe® port may hang during a Hot-plug removal event, and will not retrain until after the entire affected P-Link or G-Link is reset.

Potential Effect on System

A PCIe port may hang.

Suggested Workaround

None

Fix Planned

No fix planned

1391 Core::X86::Msr::IBS_OP_DATA3[IbsDcMissLat] May Overcount

Description

Core::X86::Msr::IBS_OP_DATA3[IbsDcMissLat] may overcount.

Core::X86::Msr::IBS_OP_DATA3[IbsDcMissLat] has overcounted if it is larger than $(\text{Core::X86::Msr::IBS_OP_DATA[IbsTagToRetCtr]} - \text{Core::X86::Msr::IBS_OP_DATA[IbsCompToRetCtr]})$.

Potential Effect on System

Performance monitoring or profiling software may produce unexpected results.

Suggested Workaround

None. Use the minimum of Core::X86::Msr::IBS_OP_DATA3[IbsDcMissLat] and $(\text{Core::X86::Msr::IBS_OP_DATA[IbsTagToRetCtr]} - \text{Core::X86::Msr::IBS_OP_DATA[IbsCompToRetCtr]})$ to approximate DC miss latency.

Fix Planned

No fix planned

1394 Processor May Cause Unexpected Collisions on SMBUS (System Management Bus)

Description

The processor may violate the Data Setup Time parameter of the SMBUS specification which may cause unexpected collisions to be observed on the SMBUS. The processor may require up to 500ns of Data Setup Time on SMBUS.

Potential Effect on System

Unexpected collisions may be observed on the SMBUS if Data Setup Time is less than 500ns. In rare cases, these collisions may prevent reading the DIMM SPD ROMs correctly.

Suggested Workaround

None

Fix Planned

No fix planned

1407 Processor May Signal Unexpected Fatal IF MCA Error

Description

Under a highly specific and detailed set of internal timing conditions, the processor may signal an unexpected fatal IF MCA error logged with MCA_STATUS_IF[ExtErrorCode] = 14d.

The resulting value in MCA_SYND_IF[ErrorInformation] will be one of the following: 22h, 23h, 2Fh, or 30h.

Potential Effect on System

System may reboot.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1415 Processor Undergoing C-State Change May Signal Unexpected Fatal IF MCA Error

Description

Under a highly specific and detailed set of internal timing conditions, the processor may signal an unexpected fatal IF MCA error logged with MCA_STATUS_IF[ExtErrorCode] = 14d when the processor is undergoing a C-State change.

The resulting value in MCA_SYND_IF[ErrorInformation] will be 25h.

Potential Effect on System

System may reboot.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1433 Processor May Log Unexpected LS MCE Error

Description

Under a highly specific and detailed set of internal timing conditions, the processor may log an uncorrectable LS MCE error of the type "HWA" (MCA::LS::MCA_STATUS_LS[ErrorCodeExt] = 0x16) with the value 0x5d00008e logged in MCA::LS::MCA_SYND_LS.

Potential Effect on System

System may log a fatal error and reset.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1450 Processor May Signal a Fatal LS MCA Error

Description

Under a highly specific and detailed set of internal timing conditions, the processor may signal a fatal LS MCA error, logged with:

- MCA_STATUS_LS[ErrorCodeExt] = 0x16 and
- MCA_SYND_LS[ErrorInformation] = 0x8e.

Potential Effect on System

System may hang or reset.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1455 PCID-Based INVLPGB May Fail to Flush Global Translations

Description

An INVLPGB instruction will fail to flush global translations on a target processor under the following conditions:

- The rAX operand has both bit 1 (valid PCID) and bit 3 (include Global) set, and
- The target processor is not currently running with the PCID set to the value encoded in EDX[27:16].

Potential Effect on System

None expected. Software is not expected to use INVLPGB with PCID-based flushing.

Suggested Workaround

None

Fix Planned

No fix planned

1464 32-Byte Misaligned Supervisor Shadow Stack Pointer or Supervisor Shadow Stack in Non-WB Memory May Result in a Non-Restartable Guest

Description

Under the either of following conditions, the processor may VMEXIT a guest with the Busy Bit set in the Supervisor Shadow Stack Token while not having completed the far transfer in the guest:

- The 32 bytes of data accessed at the new supervisor shadow stack pointer are not entirely contained within one page, or
- The memory for the new supervisor shadow stack is not mapped as WB Memory.

Potential Effect on System

None expected.

Software is not expected to create either of the conditions listed above. If the conditions occurred, the guest would not be able to be restarted.

Suggested Workaround

None

Fix Planned

No fix planned

1467 Unexpected #PF for Hypervisor Write to 2M or 1G Page When SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) Enabled

Description

A hypervisor may take an unexpected #PF Exception under the following conditions:

- SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) is enabled.
- The hypervisor writes to a 2M or larger page, and
- An SVM data structure (VMCB, VMSA, or AVIC backing page) associated with a currently running guest is allocated starting at a 2M boundary within the 2M or larger page used for the access.

Potential Effect on System

Unexpected #PF Exception

Suggested Workaround

Software should:

- Use a 4K page size for memory access within a region where an SVM page structure is allocated on a 2M-aligned boundary, or
- Avoid allocating SVM pages on 2M-aligned boundaries.

Fix Planned

No fix planned

1476 Certain PCIe® Loads May Cause PCIe Completion Timeout

Description

Under a highly specific and detailed set of internal timing conditions with certain PCIe® load conditions, a memory read may be delayed for long enough to cause a PCIe completion timeout logged at the root port, a switch, or an endpoint.

Potential Effect on System

PCIe Completion Timeout

Suggested Workaround

None recommended for systems that are not experiencing PCIe completion timeouts. If PCIe completion timeouts are encountered program the following bits to mitigate the issue:

- IOMMUL1PCIE0x00000110[1] to 1b
- IOMMUL1PCIE1x00000110[1] to 1b
- IOMMUL1PCIE2x00000110[1] to 1b
- IOMMUL1PCIE3x00000110[1] to 1b
- IOMMUL1PCIE4x00000110[1] to 1b
- IOMMUL1PCIE5x00000110[1] to 1b
- IOMMUL1PCIE6x00000110[1] to 1b
- IOMMUL1PCIE7x00000110[1] to 1b

If PCIe completion timeouts are still encountered after programming the bits mentioned above, then program the following bits to mitigate the issue:

- NBIO0PCICORE0x00000008[9] to 1b
- NBIO1PCICORE0x00000008[9] to 1b
- NBIO2PCICORE0x00000008[9] to 1b
- NBIO3PCICORE0x00000008[9] to 1b
- NBIO0PCICORE1x00000008[9] to 1b
- NBIO1PCICORE1x00000008[9] to 1b
- NBIO2PCICORE1x00000008[9] to 1b
- NBIO3PCICORE1x00000008[9] to 1b

Fix Planned

No fix planned

1488 Translation Agent May Fail to Update Dirty Bit in Page Table Entry for Address Translation Request

Description

The translation agent will

- fail to update the Dirty bit in the Page Table Entry and
- incorrectly return IW=1 in the ATS response

if all of the following conditions are met:

- The Address Translation Request (ATS) request is issued from device with No Write Flag (NW) bit set.
- The Memory requested through the ATS request has write permission.
- SEV-SNP (Secure Encrypted Virtualization - Secure Nested Paging) is disabled.

Potential Effect on System

If a page that should be marked dirty is paged out without the translation agent setting the D-bit, software will incorrectly revert back to the stored page, resulting in unpredictable system behavior.

Suggested Workaround

System software may contain the workaround for this erratum.

When the workaround is implemented, the translation agent may incorrectly set the D-bit in a page table entry for an Address Translation Request (ATS) request issued from Device with No Write Flag (NW) bit set and return IW=1 in the ATS response. This changes the behavior described line 5 of the table titled "AMD64 Access Privilege Conversion Table for ATS Request" in AMD's IOMMU specification. This behavior is not expected to cause any functional issue. However, the affected pages may be unnecessarily written out to the pagefile by software.

Fix Planned

No fix planned

1489 IOMMU Unpinned Mode Is Not Supported When Secure Address Translation Service Is Enabled

Description

When Secure Address Translation Service (SATS) support is enabled (bit 31 of the IOMMU Extended Feature Register MMIO Offset 0030h is programmed to 1b), IOMMU unpinned mode (bit 187 or HPTMode of the IOMMU Device Table Entry is programmed to 0b) is not supported.

Potential Effect on System

Unpredictable device behavior

Suggested Workaround

None. Use pinned host page tables when SATS is enabled.

Fix Planned

No fix planned

Documentation Support

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

- *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: General-Purpose and System Instructions*, order # 24594
- *AMD64 Architecture Programmer's Manual Volume 4: 128-Bit and 256-Bit Media Instructions*, order # 26568
- *AMD64 Architecture Programmer's Manual Volume 5: 64-Bit Media and x87 Floating-Point Instructions*, order # 26569
- *AMD I/O Virtualization Technology (IOMMU) Specification*, order # 48882
- *Processor Programming Reference (PPR) for AMD Family 19h Models 00h-0Fh Processors*, order # 55898

See the AMD Web site at www.amd.com for the latest updates to documents.