



Revision Guide for AMD Family 17h Models 30h-3Fh Processors

Publication # 56323	Revision: 0.78
Issue Date: March 2021	

Advanced Micro Devices 

© 2018-2021 Advanced Micro Devices, Inc. All rights reserved.

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 17h Models 30h-3Fh 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.....	9
Table 6. Cross-Reference of Processor Revision to Errata.....	11
Table 7. Cross-Reference of Errata to Package Type.....	13
Table 8. Cross-Reference of Errata to Processor Segments.....	15

Revision History

Date	Revision	Description
March 2021	0.78	Added errata #1146, #1169, #1200, #1212, #1215, #1216, #1218, #1225, #1235, #1277, #1286, #1288, #1290, #1291, #1294, #1305, #1308, and #1315.
September 2019	0.74	Initial public release.

Overview

The purpose of the *Revision Guide for AMD Family 17h Models 30h-3Fh* 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™ 7002 Series Processors

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 17 Model 30h-3Fh Processors*, order# 55803.

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 17h Models 30h-3Fh SP3 Processor Revisions

CPUID Fn0000_0001_EAX	AMD EPYC™ 7002 Series Processors
00830F10h (Rome-B0)	X

Mixed Processor Revision Support

AMD Family 17h 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.
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
00830F10h (Rome-B0)	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	CPUID Fn0000_0001 _EAX
		00830F10h (Rome-B0)
1140	Response Transaction May be Lost on Data Fabric	No fix planned
1146	PCIe® DPC (Downstream Port Containment) RP PIO (Root Port Programmed I/O) Error Reporting May Not Function Correctly	No fix planned
1150	PCIe® Incorrectly Updates Some AER (Advanced Error Reporting) Registers After Egress Blocking Error	No fix planned
1154	MOV SS Instructions May Take Multiple Breakpoints	No fix planned
1155	DMA or Peer-to-peer Accesses Using Guest Physical Addresses (GPAs) May Cause IOMMU Target Abort	No fix planned
1157	PCIe® Link Status May be Incorrect When Root Port Autonomously Changes to Gen4 Speed	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
1163	Some MCA_MISC0 Bits May Fail to Persist Through Warm Reset	No fix planned
1165	The PCIe® Link May Accumulate Correctable Errors in Some Gen3 Mode Configurations	No fix planned
1166	The PCIe® Link May Accumulate Correctable Errors in Gen4 Link Width x2 Mode	No fix planned
1169	PCIe® Error Masking May Fail to Mask Errors	No fix planned
1171	Requester ID May Be Set Incorrectly on Outbound PCIe® VDMs (Vendor Defined Messages)	No fix planned
1183	The Processor May Hang If it Receives INIT# While Already In the INIT# State	No fix planned
1185	PCIe® Receive Buffer Location May Be Incorrectly Overwritten	No fix planned
1200	xHCI Host May Hang If Full Speed or High Speed USB Hub is Connected	No fix planned
1212	IBS (Instruction Based Sampling) Micro-Op Retire Counters May Be Inaccurate	No fix planned
1215	IBS (Instruction Based Sampling) Counter Valid Value May be Incorrect After Exit From Core C6 (CC6) State	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
1235	Guest With AVIC (Advanced Virtual Interrupt Controller) Enabled May Fail to Process IPI (Inter-Processor Interrupt) Until Guest Is Re-Scheduled	No fix planned
1277	IOMMU May Mishandle Fault on Skipped Page Directory Entry Levels	No fix planned
1286	Spurious #GP May Occur When Hypervisor Running on Another Hypervisor	No fix planned
1288	A VDPPS Instruction May Fail to Record a Masked Exception in the MXCSR Register	No fix planned
1290	GMI (Global Memory Interface) Link May Hang Due to Failure To Retrain After Encountering CRC Errors	No fix planned

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

No.	Errata Description	CPUID Fn0000_0001 _EAX
		00830F10h (Rome-B0)
1291	Certain Performance Counters For Retire Based Events May Overcount	No fix planned
1294	xHCI Controller May Drop Data of an Isochronous TD (Transfer Descriptor) During Isochronous Transfer	No fix planned
1305	AHCI Controller Ignores COMINIT During HP6: HR_AwaitAlign State	No fix planned
1308	Guests With AVIC (Advanced Virtual Interrupt Controller) Enabled May Not Be Able to Program Interrupt Controller	No fix planned
1315	Two Processor System Configured With 3-link xGMI and Preferred IO Mode May Hang or Reset	No fix planned

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
1140	X
1146	X
1150	X
1154	X
1155	X
1157	X
1159	X
1160	X
1163	X
1165	X
1166	X
1169	X
1171	X
1183	X
1185	X
1200	X
1212	X
1215	X
1216	X
1218	X
1225	X
1235	X
1277	X
1286	X
1288	X
1290	X
1291	X
1294	X
1305	X

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

Errata	Package
	SP3
1308	X
1315	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
	AMD EPYC™ 7002 Series Processors
1140	X
1146	X
1150	X
1154	X
1157	X
1159	X
1160	X
1163	X
1165	X
1166	X
1169	X
1171	X
1183	X
1185	X
1200	X
1212	X
1215	X
1216	X
1218	X
1225	X
1235	X
1277	X
1286	X
1288	X
1290	X
1291	X
1294	X
1305	X
1308	X
1315	X

1140 Response Transaction May be Lost on Data Fabric

Description

Under a highly specific and detailed set of internal timing conditions, the CAKE (Data Fabric Coherent AMD socKet Extender) may lose a response transaction.

Potential Effect on System

System may hang or reset

Suggested Workaround

Program

- D18F3x2A4[28:25] to 4h, and
- D18F3x2A4[11:8] to 4h

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1146 PCIe® DPC (Downstream Port Containment) RP PIO (Root Port Programmed I/O) Error Reporting May Not Function Correctly

Description

The system resets to a state where an unmasked RP PIO error may produce unpredictable system behavior.

The DPC Extended Capability RP PIO SysError (PCIERCCFG::PCIE_RP_PIO_SYSERROR) register incorrectly inverts the sense of its enable bits, so that a bit value of 0b enables the function and a bit value of 1b disables the function. In addition, when enabled an RP PIO System Error may not be handled correctly by the processor resulting in unpredictable system behavior.

The reset value of PCIERCCFG::PCIE_RP_PIO_SYSERROR is 0000_0000h which, due to the inversion of the bits, enables RP PIO errors to be reported as a System Error. While the value of PCIERCCFG::PCIE_RP_PIO_SYSERROR is 0000_0000h, clearing any of the mask bits in PCIERCCFG::PCIE_RP_PIO_MASK will allow the corresponding error to be reported as a System Error.

Potential Effect on System

Unpredictable system behavior.

Suggested Workaround

Program PCIERCCFG::PCIE_RP_PIO_SYSERROR[18:0] to 7_0707h. This disables the DPC RP PIO System Error feature.

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1150 PCIe® Incorrectly Updates Some AER (Advanced Error Reporting) Registers After Egress Blocking Error

Description

When a PCIe egress blocking error occurs, PCIe will incorrectly update the following AER registers:

- Header Log Register
- TLP Prefix Log Register
- Root Error Status register, Multiple ERR_FATAL/NONFATAL Received bit.

Potential Effect on System

After a PCIe egress blocking error, the system will not have access to accurate diagnostic information contained in the affected registers.

Suggested Workaround

None

Fix Planned

No fix planned

1154 MOV SS Instructions May Take Multiple Breakpoints

Description

Under a highly specific and detailed set of internal timing conditions, the MOV SS instruction may prematurely clear the RFLAGS RF bit if:

- RFLAGS TF bit is asserted, and
- the instruction hits a debug breakpoint, and
- an interrupt is taken by the affected instruction before it completes.

Potential Effect on System

The affected instruction may take multiple instruction breakpoints.

Suggested Workaround

None

Fix Planned

No fix planned

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

1157 PCIe® Link Status May be Incorrect When Root Port Autonomously Changes to Gen4 Speed

Description

When the PCIe® root port autonomously changes to Gen4 speed, the following will occur:

- Link Status Register (Offset 12h) [15] Link Autonomous Bandwidth Status will incorrectly not be asserted. If Link Status Control (Offset 10h) [15] Link Autonomous Bandwidth Interrupt Enable is asserted, an interrupt may fail to occur.
- Link Status Register (Offset 12h) [14] Link Bandwidth Management Status will be incorrectly asserted. If Link Status Control (Offset 10h) [14] Link Bandwidth Management Interrupt Enable is asserted, a spurious interrupt will occur.

Potential Effect on System

Software may incorrectly handle an autonomous change to PCIe Gen4 Speed.

Suggested Workaround

None

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::::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::::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::::MCA_CTL_MASK_L3[5] or MCA::::MCA_CTL_MASK_L3[6] to 1b.

Fix Planned

No fix planned

1163 Some MCA_MISC0 Bits May Fail to Persist Through Warm Reset

Description

The following warm-reset persistent bits may incorrectly be cleared during a warm reset:

- MCA_MISC0_DE[43:32], ErrCnt
- MCA_MISC0_DE[48], Ovrflw
- MCA_MISC0_DE[50:49], ThresholdIntType
- MCA_MISC0_EX[43:32], ErrCnt
- MCA_MISC0_EX[48], Ovrflw
- MCA_MISC0_EX[50:49], ThresholdIntType
- MCA_MISC0_IF[43:32], ErrCnt
- MCA_MISC0_IF[48], Ovrflw
- MCA_MISC0_IF[50:49], ThresholdIntType

Potential Effect on System

A corrected error count in some MCA banks will be lost over a warm reset.

Suggested Workaround

None

Fix Planned

No fix planned

1165 The PCIe® Link May Accumulate Correctable Errors in Some Gen3 Mode Configurations

Description

The PCIe® link may accumulate correctable receiver errors due to spurious entries into recovery or accumulate correctable errors due to NAKs (Negative Acknowledgements) in the following configurations:

- Gen3 mode, Link width x1
- Gen3 mode, Link width x4

A link that degrades to one of the affected modes from a different mode will accumulate correctable errors as if it was originally configured to operate in the affected mode.

Potential Effect on System

Unexpected NAKs may occur, and the processor root complex may report correctable errors.

Suggested Workaround

System software may contain the workaround for this erratum.

Fix Planned

No fix planned

1166 The PCIe® Link May Accumulate Correctable Errors in Gen4 Link Width x2 Mode

Description

The PCIe® link may accumulate correctable receiver errors due to spurious entries into recovery or accumulate correctable errors due to NAKs (Negative Acknowledgements) in Gen4 Link width x2 mode.

A link that degrades to the affected mode from a different mode will accumulate correctable errors as if it was originally configured to operate in the affected mode.

Potential Effect on System

Unexpected NAKs may occur, and the processor root complex may report correctable errors.

Suggested Workaround

None

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

1171 Requester ID May Be Set Incorrectly on Outbound PCIe® VDMs (Vendor Defined Messages)

Description

The Requester ID may be set incorrectly on Outbound PCIe® VDMs (Vendor Defined Messages) if the Type field indicates:

- Broadcast from the Root Complex,
- Routed to the Root Complex, or
- Routed by ID.

VDMs will be transmitted correctly within PCIe devices attached to the following groups of processor pins, but not between the groups.

- Group 0: P0_*, G0_*, WAFL_*
- Group 1: P1_*, G1_*
- Group 2: P2_*, G2_*
- Group 3: P3_*, G3_*

MCTP (Management Component Transport Protocol) is the only defined and validated usage of PCIe VDMs on AMD platforms.

Potential Effect on System

System may be unable to route responses of VDM protocols. The receiving device may be unable to determine the source of the request.

Suggested Workaround

None

Fix Planned

No fix planned

1183 The Processor May Hang If it Receives INIT# While Already In the INIT# State

Description

The processor may hang if it receives INIT# while already in the INIT# state.

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

1185 PCIe® Receive Buffer Location May Be Incorrectly Overwritten

Description

Under a highly specific and detailed set of internal timing conditions, a PCIe® receive buffer SRAM location in the host controller may be incorrectly overwritten.

Potential Effect on System

None with BIOS incorporating AGESA RomePI-SP3_1.0.0.3 or later.

Running software prior to AGESA RomePI-SP3_1.0.0.3 may result in unpredictable system behavior and possible logging of uncorrectable parity error(s) in MCA_STATUS_NBIO due to incorrect data in PCIe buffer.

Suggested Workaround

BIOS incorporating AGESA RomePI-SP3_1.0.0.3 or later contains a workaround for this erratum.

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

1212 IBS (Instruction Based Sampling) Micro-Op Retire Counters May Be Inaccurate

Description

Under a highly specific and detailed set of internal timing conditions the processor may experience sampling inaccuracies in the following IBS (Instruction Based Sampling) counters:

- MSRC001_1035[IbsTagToRetCtr] (micro-op tag to retire count)
- MSRC001_1035[IbsCompToRetCtr] (micro-op completion to retire count)

Potential Effect on System

Software reading the IBS OP Data register may experience inaccuracies.

Suggested Workaround

None

Fix Planned

No fix planned

1215 IBS (Instruction Based Sampling) Counter Valid Value May be Incorrect After Exit From Core C6 (CC6) State

Description

If a core's IBS feature is enabled and configured to generate an interrupt, including NMI (Non-Maskable Interrupt), and the IBS counter overflows during the entry into the Core C6 (CC6) state, the interrupt may be issued, but an invalid value of the valid bit may be restored when the core exits CC6.

Potential Effect on System

The operating system may receive interrupts due to an IBS counter event, including NMI, and not observe a valid IBS register. Console messages indicating "NMI received for unknown reason" have been observed on Linux systems.

Suggested Workaround

None

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

1235 Guest With AVIC (Advanced Virtual Interrupt Controller) Enabled May Fail to Process IPI (Inter-Processor Interrupt) Until Guest Is Re-Scheduled

Description

Under a highly specific and detailed set of internal timing conditions, if a guest with AVIC enabled is about to be de-scheduled by the hypervisor, it may fail to process an IPI until after the guest is re-scheduled.

Potential Effect on System

A guest may not process an IPI until the guest is re-scheduled.

Suggested Workaround

None. Do not enable AVIC.

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

1286 Spurious #GP May Occur When Hypervisor Running on Another Hypervisor

Description

The processor may incorrectly generate a #GP fault if a hypervisor running on a hypervisor attempts to access the following secure memory areas:

- The reserved memory address region starting at FFFD_0000_0000h and extending up to FFFF_FFFF_FFFFh.
- ASEG and TSEG memory regions for SMM (System Management Mode)
- MMIO APIC Space

Potential Effect on System

Software running a hypervisor on a hypervisor may encounter an unexpected #GP fault.

Suggested Workaround

If CPUID bit fn8000_000A EDX[28] = 0b, then:

- Hypervisor software can trap #GP faults that potentially have this issue and ignore #GP faults that were erroneously generated.

If CPUID bit fn8000_000A EDX[28] = 1b, then the issue has been fixed and no workaround is necessary.

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

1290 GMI (Global Memory Interface) Link May Hang Due to Failure To Retrain After Encountering CRC Errors

Description

GMI Link may hang due to failure to retrain after encountering correctable CRC errors. As a result, the following system errors may be observed:

- EX Watchdog Timeout (MCA::EX::MCA_STATUS_EX[ErrorCodeExt] = 0h) on one or more cores
- PCIe Completion Timeout (PCIe Uncorrectable Error Status register bit 14 = 1b)
- DF HW_ASSERT error (MCA::PIE::MCA_STATUS_PIE[ErrorCodeExt] = 0h)

Any of the above errors may be preceded by one or more occurrences of the following error:

- GMI correctable CRC errors (MCA::PIE::MCA_STATUS_PIE[ErrorCodeExt] = 2h, MCA::PIE::MCA_SYND_PIE[17] = 1b)

Potential Effect on System

System may hang or reset.

Suggested Workaround

System software may contain a workaround for this issue.

Fix Planned

No fix planned

1291 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]
- 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

None

Software may enable another counter by setting PMC0x22[4] with a value of 1b. After reading the initial counter, if the second counter (with PMC0x22[4] set to 1b) is read and the count is zero, then the overcounting did not occur.

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

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

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

1315 Two Processor System Configured With 3-link xGMI and Preferred IO Mode May Hang or Reset

Description

A two processor system configured with 3-link xGMI and preferred IO mode may hang or reset.

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

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 17h Models 30h-3Fh Processors*, order# 55803

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