POWER9 Processor SCM

Hardware Errata Notice

DD 2.2

OpenPOWER

Version 1.2
19 April 2019—IBM Confidential

Printed in the United States of America April 2019

IBM Systems
294 Route 100, Building SOM4
Somers, NY 10589-3216

The IBM home page can be found at **ibm.com®**.

Version 1.2
19 April 2019—IBM Confidential

Introduction

For the design revision levels specified (DD 2.2), this document identifies differences between the actual operation of the IBM® POWER9™ processor SCM devices and the system design described in the POWER9 User Manual and Datasheets. Differences that are not visible to the user are not described.

Notes:

- Some of the workarounds described in this document might involve the use of IBM modules to affect, or to reduce the effect of, the erratum. These modules are typically supplied with the POWER9 devices. If, however, you do not have the module required for a workaround, contact IBM to acquire it.
- This is a working document. Check regularly with your IBM technical representative to verify that you have the most current version. For more information, contact OpenPOWER@us.ibm.com.

Revision Levels Covered

This document includes information about errata that apply to the following design revision levels:

- POWER9 processor SCM revision DD 2.2 and later

This document does not contain information about POWER9 processor levels before DD 2.2. Any statements in this document about other revision levels of these products are for reference only. For errata information about these revision levels, see the current Errata Notice for the specific revision of the product.

Errata Categories

Each erratum is assigned to a category numbered 1 - 5 based on its impact on system performance and the ability of any proposed workaround to minimize that impact. *Table 1* explains the meaning of each category.

Table 1. Errata Categories

Errata Category	Definition
1	Major impact; no workaround available. An erratum is said to have a major impact if it results in a system crash, a hard failure, an unrecoverable soft failure, significant performance degradation, or the storage of incorrect data.
2	Major impact; workaround is impractical to implement, or a substantial risk exists of encountering the same problem or additional problems, including performance issues, after the workaround is implemented.
3	Major impact; workaround available. Application of the workaround either eliminates the problem, or reduces it to a minor impact issue.
4	Minor impact; no workaround available. An erratum is said to have a minor impact if it results in slight-to-moderate performance degradation or is a functional variance from the specification.
5	Minor impact; workaround is available. An erratum is said to have a minor impact if it results in slight-to-moderate performance degradation or is a functional variance from the specification. Application of the workaround eliminates the problem.

Terminology

The following section headings and terms are used in this document:

Erratum #n	The errata are numbered sequentially.
Erratum	An erratum is any difference between the actual operation of the specific design revision level and the design specification.
Abstract	The abstract is a brief description of the erratum.
CQ ID	Common queue ID.
Impact	A category is assigned to an erratum based on its impact on system performance and the ability of any proposed workaround to minimize that impact. See *Table 1* on page 3.
Published Date	The date that the erratum was first published.
Applies to	Indicates which devices and revision levels are affected by the erratum.
Status	This is the current status of the erratum discussion and the current plan for the future.
Workarounds	This section lists any actions that customers can take to reduce or eliminate the effects of the erratum. These actions supplement any workarounds embedded in the IBM code. If you are not using the operating system or firmware supplied or recommended by IBM, contact OpenPOWER@us.ibm.com.

Summary of Errata

Table 2 summarizes the errata described in this document.

Table 2. Summary of Errata

Erratum Number	Abstract	Errata Category	Workaround Available?	See Page
1	The **DAWR** feature is disabled on DD2.2.	2	No	5
2	Possible system hang can occur.	5	Yes	5
3	Problem with the **mftexasr** instruction in TM suspend mode.	4	No	6
4	Disable_pb_tce_arbitration of the pbcqhwcfg Register must be set to '1'.	4	Yes	6
5	Enhanced peer-to-peer mode is not supported.	4	Yes	7
6	The **armwf*** commands only work on cache-line addresses. The **armwf_dec_b** is not supported	4	Yes	7
7	All PMU events that measure reload latency (PM_MRK_DATA_FROM_*_CYC) overcount the latency of marked noncacheable loads.	4	No	8

Erratum #1: DAWR feature is disabled on DD2.2

Abstract:	The **DAWR** feature is disabled on DD2.2.
CQ ID	None.
Impact:	Category 2. Major impact; workaround is impractical to implement, or a substantial risk exists of encountering the same problem or additional problems, including performance issues, after the workaround is implemented. The **DAWR** feature does not work on DD2.2.
Published Date:	12 February 2018.
Applies To:	POWER9 processor SCM.
Status:	Permanent errata.
Workaround	None.

Erratum #2: System hang can occur

Abstract:	Possible system hang can occur. Conditions exist where a thread can starve out a neighbor thread.
CQ ID	None.
Impact:	Category 5. Minor impact; workaround is available. An erratum is said to have a minor impact if it results in slight-to-moderate performance degradation or is a functional variance from the specification. Application of the workaround eliminates the problem.
Published Date:	12 February 2018.
Applies To:	POWER9 processor SCM.
Status:	Permanent errata.
Workaround:	The HDEC should not be set for longer than 5 seconds. Taking a decrementer interrupt prevents the hang condition.

Erratum #3: The mftexasr instruction in TM suspend mode

Abstract:	Problem with the **mftexasr** instruction in TM. When the **mftexasr** instruction is executed in suspend mode unpredictable data can be returned.
CQ ID	None.
Impact:	Category 4. Minor impact; no workaround available. An erratum is said to have a minor impact if it results in slight-to-moderate performance degradation or is a functional variance from the specification. Unpredictable data can be returned when the **mftexasr** instruction is executed in suspend mode.
Published Date:	12 February 2018.
Applies To:	POWER9 processor SCM.
Status:	Permanent errata.
Workaround:	None.

Erratum #4: Bit disable_pb_tce_arbitration Must Always be On

Abstract:	The disable_pb_tce_arbitration bit must always be on.
CQ ID	HW368162.
Impact:	Category 4. Minor impact; a workaround is available. These are less severe problems in that they might occur less frequently or have fewer side effects. There is a workaround available for the problem. The workaround might or might not be sufficient as a permanent fix.
Published Date:	05 April 2017.
Applies to:	POWER9 processor SCM.
Status:	No fix is planned.
Workaround:	This mode should be set to ‘1’ in the register PBCQHWCFG, bit 60. Change to the initialization sequence.

Erratum #5: Enhanced peer-to-peer mode not supported

Abstract:	Enhanced peer-to-peer (P2P) mode is not supported.
CQ ID	HW400810.
Impact:	Category 4. Minor impact; a workaround is available. These are less severe problems in that they might occur less frequently or have fewer side effects. There is a workaround available for the problem. The workaround might or might not be sufficient as a permanent fix.
Published Date:	05 April 2017.
Applies to:	POWER9 processor SCM.
Status:	No fix is planned.
Workaround:	Use normal peer-to-peer mode. Documentation of enhanced peer-to-peer mode has been removed.

Erratum #6: ARMWF* commands only work on cache-line address boundaries

Abstract:	The **armwf*** commands only work on cache-line address boundaries.
CQ ID	HW403589.
Impact:	Category 4. Minor impact; a workaround is available. These are less severe problems in that they might occur less frequently or have fewer side effects. There is a workaround available for the problem. The workaround might or might not be sufficient as a permanent fix.
Published Date:	05 April 2017.
Applies to:	POWER9 processor SCM.
Status:	No fix is planned.
Workaround:	Only use **armwf*** commands on cache-line boundaries. Because of this restriction, it is not possible to use the **armwf_dec_b** command.

Erratum #7: All PMU events that measure reload latency overcount latency of marked noncacheable loads

Abstract: All PMU events that measure reload latency (PM_MRK_DATA_FROM_*_CYC) overcount the latency of marked noncacheable loads.

Impact: Category 4. Minor impact; no workaround available. An erratum is said to have a minor impact if it results in slight-to-moderate performance degradation or is a functional variance from the specification.

Published Date: 19 April 2019

Applies To: POWER9 processor SCM (all releases)

Status: A fix is planned for the POWER10 processor.

Description

Marked reload-source cycle events are intended to count cycles from the time a marked L1 miss is reported until the time the corresponding reload arrives at the core. An internal latch is set when the marked L1 miss occurs, and it is cleared when the reload occurs. However, for noncacheable loads, this latch is not cleared at reload time. The latch remains set until the next marked reload occurs. This results in overcounting for the following events whenever noncacheable loads are eligible for marking with either random instruction sampling or random event sampling:

PM_MRK_DATA_FROM_DL2L3_MOD_CYC
PM_MRK_DATA_FROM_DL2L3_SHR_CYC
PM_MRK_DATA_FROM_DL4_CYC
PM_MRK_DATA_FROM_DMEM_CYC
PM_MRK_DATA_FROM_L2_1_MOD_CYC
PM_MRK_DATA_FROM_L2_1_SHR_CYC
PM_MRK_DATA_FROM_L2_CYC
PM_MRK_DATA_FROM_L2_DISP_CONFLICT_LDHITST_CYC
PM_MRK_DATA_FROM_L2_DISP_CONFLICT_OTHER_CYC
PM_MRK_DATA_FROM_L2_MEPF_CYC
PM_MRK_DATA_FROM_L2MISS_CYC
PM_MRK_DATA_FROM_L2_NO_CONFLICT_CYC
PM_MRK_DATA_FROM_L3_1_ECO_MOD_CYC
PM_MRK_DATA_FROM_L3_1_ECO_SHR_CYC
PM_MRK_DATA_FROM_L3_1_MOD_CYC
PM_MRK_DATA_FROM_L3_1_SHR_CYC
PM_MRK_DATA_FROM_L3_CYC
PM_MRK_DATA_FROM_L3_DISP_CONFLICT_CYC
PM_MRK_DATA_FROM_L3_MEPF_CYC
PM_MRK_DATA_FROM_L3MISS_CYC
PM_MRK_DATA_FROM_L3_NO_CONFLICT_CYC
PM_MRK_DATA_FROM_LL4_CYC
PM_MRK_DATA_FROM_LMEM_CYC
PM_MRK_DATA_FROM_MEMORY_CYC
PM_MRK_DATA_FROM_OFF_CHIP_CACHE_CYC
PM_MRK_DATA_FROM_ON_CHIP_CACHE_CYC

PM_MRK_DATA_FROM_RL2L3_MOD_CYC
PM_MRK_DATA_FROM_RL2L3_SHR_CYC
PM_MRK_DATA_FROM_RL4_CYC
PM_MRK_DATA_FROM_RMEM_CYC
PM_MRK_LD_MISS_EXPOSED_CYC
PM_MRK_LD_MISS_L1_CYC

To estimate the reliability of these latency events, the following event can be counted for the application of interest. This event shows the total number of unmarked noncacheable loads. A high value indicates probable inaccuracies for the latency events.

- PM_LD_CACHE_INHIBITED

Workaround

No workaround is available.

Revision Log

Each release of this document supersedes all previously released versions. The revision log lists all significant changes made to the document since its initial release. In the rest of the document, change bars in the margin indicate that the adjacent text was modified from the previous release of this document.

Revision Date	Description
19 December 2018	Version 1.2. • Revised *Table 2 Summary of Errata* on page 4. • Added *Erratum #7: All PMU events that measure reload latency overcount latency of marked non-cacheable loads* on page 8.
29 June 2018	Version 1.1. • Added *Erratum #4: Bit disable_pb_tce_arbitration Must Always be On* on page 6. • Added *Erratum #5: Enhanced peer-to-peer mode not supported* on page 7. • Added *Erratum #6: ARMWF* commands only work on cache-line address boundaries* on page 7.
19 March 2018	Version 1.0 (initial release).