POWER9 Processor DD 2.1 Use Restrictions

Application Note

OpenPOWER

Version 1.0
5 June 2019

IBM ®

Printed in the United States of America June 2019

All information contained in this document is subject to change without notice. The products described in this document are NOT intended for use in applications such as implantation, life support, or other hazardous uses where malfunction could result in death, bodily injury, or catastrophic property damage. The information contained in this document does not affect or change IBM product specifications or warranties. Nothing in this document shall operate as an express or implied indemnity under the intellectual property rights of IBM or third parties. All information contained in this document was obtained in specific environments, and is presented as an illustration. The results obtained in other operating environments may vary.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED ON AN “AS IS” BASIS. IBM makes no representations or warranties, either express or implied, including but not limited to, warranties of merchantability, fitness for a particular purpose, or non-infringement, or that any practice or implementation of the IBM documentation will not infringe any third party patents, copyrights, trade secrets, or other rights. In no event will IBM be liable for damages arising directly or indirectly from any use of the information contained in this document.

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

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

Version 1.0
5 June 2019

Contents

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
5 June 2019	Version 1.0 (initial release).

Introduction

For the design revision level specified (DD 2.1), this document identifies use restrictions for the IBM® POWER9™ processor SCM devices and system design considerations.

Note: 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 use restrictions that apply to the following design revision level:

- POWER9 processor SCM revision DD 2.1

This document does not contain information about POWER9 processor levels before DD 2.1. Any statements in this document about other revision levels of these products are for reference only.

Terminology

The following terms are used in this document:

Use Restrictions
Refers to hardware or software limits on the use of a feature or function for the specified design revision level.

Limitation
A brief description of the use restriction.

Workarounds
This section lists any actions that customers can take to reduce or eliminate the effects of the use restriction. 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.

KVM
Kernel-based virtual machine. A virtual machine implementation that uses the operating system kernel (typically Linux).

MCD
Memory coherence directory.

Meltdown
The vulnerability known as Meltdown allows user-level code to infer the contents of the kernel memory.

While this vulnerability does not enable an external unauthorized party to gain access to a machine, it could allow a party with access to a system to access unauthorized data.

PMU
Performance monitor unit.

Spectre
The vulnerability known as Spectre allows user-level code to infer data from unauthorized memory.

While this vulnerability does not allow an external unauthorized party to gain access to a machine, it could allow a party with access to a system to access unauthorized data.

TM	Transactional memory.
TPM	Trusted platform module.
WOF	Workload optimized frequency.

Summary of Restrictions

Table 1 summarizes the use restrictions described in this document.

Table 1. Summary of POWER9 DD 2.1 Use Restrictions

Section	Abstract	See Page
1.1	Spectre and Meltdown Protections	9
1.2	Memory Configuration	9
1.3	Transactional Memory	10
1.4	Secure Boot	10
1.5	KVM	10
1.6	Copy-Paste	11
1.7	Idle States	11
1.8	Performance Monitoring Unit	11
1.9	OpenCAPI	11

Related Documents

The following documents can be helpful when reading this application note. Contact your IBM representative to obtain any documents that are not available through the IBM Portal for OpenPOWER, an online IBM technical library or the OpenPOWER Foundation web site.

POWER9 Processor User's Manual

POWER9 Performance Monitor Unit User's Guide

Power ISA User Instruction Set Architecture - Book I (Version 3.0B)

Power ISA Virtual Environment Architecture - Book II (Version 3.0B)

Power ISA Operating Environment Architecture - Book III (Version 3.0B)

1. POWER9 Processor DD 2.1 Use Restrictions

Consider the following use restrictions when designing with the POWER9 processor DD 2.1 SCM.

1.1 Spectre and Meltdown Protections

Limitation

No Spectre and Meltdown protections are available if building from the main OpenPOWER firmware repository. Protections exist if using OP910 firmware but offer limited control.

Description

The POWER9 processor is vulnerable to the reported security flaws exposed by speculative execution, collectively known as Spectre and Meltdown. With OP910, the supported security mitigations default to Kernel only protection [initialization (init) level 0] and cannot be disabled. Additional user-to-user and kernel-to-user protections are available with init level 1. No mitigations exist for DD 2.1 outside of OP910.

Workaround

None

1.2 Memory Configuration

Limitation

The maximum memory configuration is restricted to 512 GB per SCM if using OP910 firmware.

Description

To resolve an errata with the memory coherence directory (MCD), the memory mapping on the POWER9 DD 2.1 chip was modified to accommodate GPU memory. As a result of the workaround, the new mapping limits the maximum memory group size to 512 GB per socket.

Workaround

This restriction can be removed by cherry-picking[1] the following commits to disable the MCD.

- Commit 1
- Commit 2

[1] https://git-scm.com/docs/git-cherry-pick

1.3 Transactional Memory

Limitation	Not supported.
Description	On the POWER9 DD 2.1 processor, transactional memory (TM) is configured by the firmware to always abort a transaction when a TM suspend occurs. Therefore, `tsuspend` causes a transaction to be aborted and rolled back. Kernel exceptions also cause the transaction to be aborted and rolled back and the exception does not occur. If the userspace constructs a sigcontext that enables a TM suspend, the sigcontext will be rejected by the kernel. This mode is advertised to users with HWCAP2[PPC_FEATURE2_HTM_NO_SUSPEND] set. HWCAP2[PPC_FEATURE2_HTM] is not set in this mode.
Workaround	None.

1.4 Secure Boot

Limitation	No secure boot or trusted platform module (TPM) support if using OP910 firmware.
Description	Secure boot and TPM are an integral part of platform security. Though OP910 was the target of several security fixes orthogonal to secure boot, it is missing key components to enable comprehensive protection of platform boot code.
Workaround	Support available starting with OpenPOWER firmware tag v2.0.

1.5 KVM

Limitation	Not supported.
Description	KVM support is limited to POWER9 offerings with the POWER9 processor (DD 2.2 and above) through the Ubuntu 18.04 and Red Hat Enterprise Linux 7.5 and 8.0 distributions for IBM Power LE (POWER9).
Workaround	None.

1.6 Copy-Paste

Limitation No user-space support for copy-paste instructions.

Description The copy-paste facility introduced in the POWER9 chip provides an optimized mechanism for a user-space application to copy a cache line. Due to known errata with the POWER9 processor (DD 2.1), copy-paste support is restricted.

Workaround None.

1.7 Idle States

Limitation Only STOP 0, 1, and 2 are supported if using OP910 firmware.

Description The POWER9 architecture supports the following set of STOP states: 0, 1, 2, 4, 5, and 11. Support for STOP 4, STOP 5, and STOP 11 is missing from OP910. As a result, users might experience higher idle power and the maximum supported frequency under WOF conditions will be capped below the boost maximum.

Workaround Full idle state and WOF support is available starting with OpenPOWER firmware tag v2.0.

1.8 Performance Monitoring Unit

Limitation Limited performance monitoring unit (PMU) support.

Description Due to known errata, some PMU events are disabled. As such, performance monitoring tools that depend on the PMU might not function as expected.

Workaround Contact your IBM technical representative for details.

1.9 OpenCAPI

Limitation Not supported.

Description Due to known errata, the 25G PHY does not support running at the full production speed of 25.78125 GHz. As such, OpenCAPI devices must only run at lower speed for development purposes.

Workaround None.