

Intel® Xeon Phi™ Coprocessor x100 Product Family

Specification Update

March 2015



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: <http://www.intel.com/design/literature.htm>


Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at <http://www.intel.com/> or from the OEM or retailer.

No computer system can be absolutely secure.

Intel, Xeon, Xeon Phi, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2015, Intel Corporation. All rights reserved.




Contents

Revision History	4
1.0 Preface	5
2.0 Identification Information	8
3.0 Summary Table of Changes	17
4.0 Errata	20
5.0 Specification Changes	32
6.0 Specification Clarifications	33
7.0 Documentation Changes	34



Revision History

Version	Description	Date
2.0	<ul style="list-style-type: none"> Initial release 	November 2012
2.1	<ul style="list-style-type: none"> Revised Table 2-3 Revised Table 2-4 Revised Summary Table of Changes description Revised CD15, CD24, CD26 	January 2013
2.2	<ul style="list-style-type: none"> Added errata CD39-CD45 Added Specification Changes CD1 Added Mixing Coprocessors Within a Platform Revised Table 2-1 Revised Table 2-3 	February 2013
2.3	<ul style="list-style-type: none"> Added errata CD46-CD53 Revised CD38 Revised Table 2-2 Revised Table 2-3 Revised Table 2-4 	March 2013
2.4	<ul style="list-style-type: none"> Revised Table 2-3 Revised Table 2-4 	April 2013
2.5	<ul style="list-style-type: none"> Added CD54-56 Revised Table 2-2 Revised Table 2-3 Revised Table 2-4 	May 2013
2.6	<ul style="list-style-type: none"> Revised Table 2-1 Revised "Mixing Coprocessors Within a Platform" Revised errata CD22, CD23, CD27, CD40, CD56 Revised "Errata Summary Table" Revised Table 2-2 Revised Table 2-3 Revised Table 2-4 	June 2013
2.7	<ul style="list-style-type: none"> Revised Table 2-3 Revised Table 2-4 	July 2013
2.8	<ul style="list-style-type: none"> Revised Table 2-4 	September 2013
2.9	<ul style="list-style-type: none"> Revised Table 2-2 Revised Table 2-3 Revised Table 2-4 Revised erratum CD56 	December 2013
3.0	<ul style="list-style-type: none"> Revised CD31 Revised Table 2-4 	January 2014
3.1	<ul style="list-style-type: none"> Revised Table 2-4 	February 2014
3.2	<ul style="list-style-type: none"> Revised Table 2-2 Revised Table 2-4 	April 2014
3.3	<ul style="list-style-type: none"> Revised Table 2-2 	January 2015
3.4	<ul style="list-style-type: none"> Revised Table 2-2 Revised Table 2-3 Revised Table 2-4 	February 2015
3.5	<ul style="list-style-type: none"> Revised Table 2-2 Revised Table 2-4 	March 2015



1.0 Preface

This document is an update to the specifications contained in the [Affected Documents](#) table below. This document is a compilation of device and documentation errata, specification clarifications and changes. It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools for the Intel® Xeon Phi™ coprocessor.

Information types defined in [Nomenclature](#) are consolidated into the specification update and are no longer published in other documents.

The errata are separated into sub-groups to assist in understanding the status of the erratum and what action, if any, needs to be taken to address the erratum. The names and definitions of the sub-groups are detailed below.

1.1 Affected Documents

Document Title	Document Number
<i>Intel® Xeon Phi™ Coprocessor Datasheet</i>	intel.com/content/www/us/en/processors/xeon/xeon-technical-resources.html
<i>Intel® Xeon Phi™ Coprocessor System Software Developers Guide</i>	intel.com/content/www/us/en/processors/xeon/xeon-technical-resources.html
<i>Intel® Xeon Phi™ Coprocessor Instruction Set Architecture Reference Manual</i>	software.intel.com/mic-developer
<i>Intel® Manycore Platform Software Stack Readme</i>	software.intel.com/mic-developer

1.2 Related Documents

Document Title	Document Number/Location
<i>Manycore Platform Software Stack Release</i>	software.intel.com
<i>AP-485, Intel® Processor Identification and the CPUID Instruction</i>	241618



Document Title	Document Number/ Location
<i>Intel® 64 and IA-32 Architecture Software Developer's Manual, Volume 1: Basic Architecture</i>	253665
<i>Volume 2A: Instruction Set Reference, A-M</i>	253666
<i>Volume 2B: Instruction Set Reference, N-Z</i>	253667
<i>Volume 3A: System Programming Guide, Part 1</i>	253668
<i>Volume 3B: System Programming Guide, Part 2</i>	253669
<i>ACPI Specifications</i>	www.acpi.info
<i>PCIe® Base 2.1 Specification</i>	www.pcisig.com



1.3 Nomenclature

TA Number is a code used to identify different manufacturing “SKUs” of the Intel® Xeon Phi™ coprocessor as identified in a table below.

Specification Changes / Clarifications are modifications to the currently published specifications. These changes will be incorporated in the next release of the specification.

Documentation Changes include typographical errors, omissions, or incorrect information from the current published specifications. These will be incorporated in the next release of the specification.

Specification changes, specification clarifications, and documentation changes are removed from the specification update when the appropriate changes are made to the appropriate product specification or user documentation.



2.0 Identification Information

2.1 Component Identification

The Intel® Xeon Phi™ coprocessor’s silicon stepping can be identified by the following register contents.

Table 2-1. Intel® Xeon Phi™ Coprocessor Silicon Signature/Version

Reserved	Extended Family ¹	Extended Model ²	Reserved	Processor Type ³	Family Code ⁴	Model Number ⁵	Stepping ID ⁶
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	00000000b	0000b		00b	1011b	0001b	B0=0001b B1=0011b C0=0010b

Notes:

1. The Extended Family, bits [27:20] are used in conjunction with the Family Code, specified in bits [11:8], to indicate whether the processor belongs to the Intel386™, Intel486™, Pentium®, Pentium® Pro, Pentium® 4, Intel® Core™, or Intel® Many Integrated Core Architecture processor family.
2. The Extended Model, bits [19:16] in conjunction with the Model Number, specified in bits [7:4], are used to identify the model of the processor within the processor family.
3. The Processor Type, specified in bits [13:12] indicates whether the processor is an original OEM processor, an OverDrive processor, or a dual processor (capable of being used in a dual processor system).
4. The Family Code corresponds to bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register.
5. The Model Number corresponds to bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register.
6. The Stepping ID in bits [3:0] indicates the revision number of that model.

When EAX is initialized to a value of '1', the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number, and Stepping ID in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.

Cache and TLB descriptor parameters are provided in the EAX, EBX, ECX and EDX registers after the CPUID instruction is executed with a 2 in the EAX register.



2.2 Intel® Xeon Phi™ Coprocessor Identification

The Intel® Xeon Phi™ coprocessor is manufactured in the following SKUs:

Table 2-2. Intel® Xeon Phi™ Coprocessor Identification

SKU	TA#	Silicon Stepping	Revision ID	Device ID	Subsystem ID	Core Count	Core Frequency	Cache Size	Memory Size	Memory Speed	Cooling Solution	TDP
SE10P	G71513-2XX	B0	0x10	0x225C	0x2500	61	1.1 GHz Turbo: n/a ¹	30.5 MB	8 GB	5.5 GT/s	Passive	300 W
SE10X	G65756-2XX	B1	0x11	0x225C	0x2500	61	1.1 GHz Turbo: n/a ¹	30.5 MB	8 GB	5.5 GT/s	None	300 W
SE10P	G65757-2XX G78930-2XX	B1	0x11	0x225C	0x2500	61	1.1 GHz Turbo: n/a ¹	30.5 MB	8 GB	5.5 GT/s	Passive	300 W
SC5110P	G65758-2XX	B1	0x11	0x2250	0x2500	60	1.053 GHz Turbo: n/a ¹	30 MB	8 GB	5 GT/s	Passive	225 W
SC5110PEB	H13742-2XX	B1	0x11	0x2250	0x2500	60	1.053 GHz Turbo: n/a ¹	30 MB	8 GB	5 GT/s	Passive	225 W
SC5110PKIT	G65758-3XX	B1	0x11	0x2250	0x2500	60	1.053 GHz Turbo: n/a ¹	30 MB	8 GB	5 GT/s	Passive	225 W
SC31S1P	G78927-2XX	B1	0x11	0x225E	0x2500	57	1.1 GHz Turbo: n/a ¹	28.5 MB	8 GB	5 GT/s	Passive	270 W
SC7120X	G65756-3XX	C0	0x20	0x225C	0x7D91	61	1.238 GHz Turbo: 1.333 GHz	30.5 MB	16 GB	5.5 GT/s	None	300 W
SC7120P	G65757-3XX	C0	0x20	0x225C	0x7D95	61	1.238 GHz Turbo: 1.333 GHz	30.5 MB	16 GB	5.5 GT/s	Passive	300 W
SC7120PEB	H13647-3XX	C0	0x20	0x225C	0x7D95	61	1.238 GHz Turbo: 1.333 GHz	30.5 MB	16 GB	5.5 GT/s	Passive	300 W
SC7120D	H16728-2XX	C0	0x20	0x225C	0x7D9D 0x7D91 ²	61	1.238 GHz Turbo: 1.333 GHz	30.5 MB	16 GB	5.5 GT/s	None	300 W



Table 2-2. Intel® Xeon Phi™ Coprocessor Identification

SKU	TA#	Silicon Stepping	Revision ID	Device ID	Subsystem ID	Core Count	Core Frequency	Cache Size	Memory Size	Memory Speed	Cooling Solution	TDP
SC7120A	G86604-3XX	C0	0x20	0x225C	0x7D99	61	1.238 GHz Turbo: 1.333 GHz	30.5 MB	16 GB	5.5 GT/s	Active	280 W
SC5120D	G61933-2XX	C0	0x20	0x2250	0xD804 0xDD94 ²	60	1.053 GHz Turbo: n/a ¹	30 MB	8 GB	5.5 GT/s	None	245 W
SC3120P	G82474-3XX	C0	0x20	0x225D	0x3604 0x3C94 ²	57	1.1 GHz Turbo: n/a ¹	28.5 MB	6 GB	5 GT/s	Passive	300 W
SC3120PEB	H13740-3XX	C0	0x20	0x225D	0x3604 0x3C94 ²	57	1.1 GHz Turbo: n/a ¹	28.5 MB	6 GB	5 GT/s	Passive	300 W
SC3120A	G65759-3XX	C0	0x20	0x225D	0x3608 0x3C98 ²	57	1.1 GHz Turbo: n/a ¹	28.5 MB	6 GB	5 GT/s	Active	300 W
SC3120AEB	H13668-3XX	C0	0x20	0x225D	0x3608 0x3C98 ²	57	1.1 GHz Turbo: n/a ¹	28.5 MB	6 GB	5 GT/s	Active	300 W
SC3120AIB	H13708-3XX	C0	0x20	0x225D	0x3608 0x3C98 ²	57	1.1 GHz Turbo: n/a ¹	28.5 MB	6 GB	5 GT/s	Active	300 W
SC3120AKIT	G65759-3XX	C0	0x20	0x225D	0x3608 0x3C98 ²	57	1.1 GHz Turbo: n/a ¹	28.5 MB	6 GB	5 GT/s	Active	300 W

Note: Coprocessors with serial numbers where the 5th, 6th, and 7th digit show less than 241 may be affected by erratum CD38.

Notes:

1. This SKU does not support Turbo Mode.
2. Contact your Intel Support Representative for more information.

2.3 Mixing Coprocessors Within a Platform

Mixing coprocessors of different steppings but the same SKU (as per Table 2-2, "Intel® Xeon Phi™ Coprocessor Identification") is supported per the following table:

Stepping	B0	B1	C0
B0	Supported	Supported	Unsupported
B1	Supported	Supported	Supported
C0	Unsupported	Supported	Supported

The only supported features are those of the coprocessor with the lowest stepping. Customers are fully responsible for the validation of their system configurations.



2.4 Firmware Revision

Each unique PCI Express card with associated coprocessor silicon stepping has a firmware image that, when applied, constitutes a supported PCI Express card (i.e., a specified PCI Express card includes a firmware revision). The proper firmware revision must be loaded on each Intel® Xeon Phi™ coprocessor in a system. The proper firmware revision is defined as the minimum firmware revision available from Intel for a given coprocessor silicon stepping. Subsequent updates to the minimum revision may be considered optional depending on added functionality. Any Intel® Xeon Phi™ coprocessor that does not have the minimum firmware revision loaded is considered to be operating out of specification. Contact your Intel Representative to receive the latest firmware revision.

Table 2-3. Intel® Xeon Phi™ Coprocessor Firmware Revision Guide

Firmware Revision	SMC Revision	Customer Release Date	Intended Stepping	Memory Speed Enabled	Workaround for Errata
375-1	1.6	November 2012	B0 B1	5.0 GT/s or 5.5 GT/s	CD11
375-5	1.7 ¹	December 2012	B0 B1	5.0 GT/s or 5.5 GT/s	
383-2	1.11 ¹	March 2013	B0 B1	5.0 GT/s or 5.5 GT/s	CD43
385-1	1.13 ¹	March 2013	B0 B1	5.0 GT/s or 5.5 GT/s	CD54 ² , CD55
386-2	1.14 ¹	May 2013	B0 B1 C0	5.0 GT/s or 5.5 GT/s	CD54
386-3	1.15 ¹	July 2013	B0 B1 C0	5.0 GT/s or 5.5 GT/s	
390-2	1.16 ¹	December 2013	B0 B1 C0	5.0 GT/s or 5.5 GT/s	
391-2	1.17 ¹	December 2014	B0 B1 C0	5.0 GT/s or 5.5 GT/s	

Note: The latest Intel® Xeon Phi™ coprocessor firmware can be found within the Intel® Manycore Platform Software Stack, which is available at <http://software.intel.com/en-us/articles/intel-manycore-platform-software-stack-mpss>.

Notes:

1. This SMC revision depends on a SMC bootloader update with a version number of 1.8, that is supplied by this Intel® MPSS release.
2. The Intel® MPSS release contains a partial fix for this erratum. A full fix has been implemented in a later Intel® MPSS release.



2.5 Intel® Manycore Platform Software Stack (MPSS) Revision

The Intel® Manycore Platform Software Stack (MPSS) is the collection of software which enables use of the Intel® Xeon Phi™ coprocessor. It is made up of several components:

- A Linux* operating system which runs on the Intel® Xeon Phi™ coprocessor. This includes a customized kernel, utility programs, and run-time libraries
- A set of host device drivers that enables both programmatic and user-level access to the Intel® Xeon Phi™ coprocessor
- Low-level device drivers and library support for high-performance data transfer between the host and the Intel® Xeon Phi™ coprocessor
- System management utilities for monitoring and administering the Intel® Xeon Phi™ coprocessor

Table 2-4. Intel® Manycore Platform Software Stack Revision Guide

Revision	Flash Revision	Customer Release Date	Supported Stepping	Supported Operating Systems	Workaround for Errata
2.1.4346-16	375-1	November 2012	B0 B1	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3 SUSE Linux Enterprise Server* 11 SP1, SP2	CD31
2.1.4892-15	375-5	December 2012	B0 B1	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3 SUSE Linux Enterprise Server* 11 SP1, SP2	
2.1.5889-14	383-2	March 2013	B0 B1	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3 SUSE Linux Enterprise Server* 11 SP1, SP2	CD31, CD35, CD42
2.1.5889-16	385-01	March 2013	B0 B1	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3 SUSE Linux Enterprise Server* 11 SP1, SP2	CD55
2.1.6720-13	386-2	May 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP1, SP2 Microsoft Windows* 7 Enterprise Microsoft Windows* 8 Enterprise Microsoft Windows* Server 2008 R2 Microsoft Windows* Server 2012	CD32
2.1.6720-15	386-3	July 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP1, SP2	
2.1.6720-16	386-3	August 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP1, SP2	
2.1.6720-19	386-3	September 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP1, SP2	
2.1.6720-21	386-3	October 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP1, SP2	
2.1.6720-21	386-3	March 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP1, SP2	



Table 2-4. Intel® Manycore Platform Software Stack Revision Guide

Revision	Flash Revision	Customer Release Date	Supported Stepping	Supported Operating Systems	Workaround for Errata
3.1	386-3	October 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012	CD56
3.1.1	390-2	December 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012	
3.1.2	390-2	January 2013	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012	
3.1.4	390-2	March 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012	
3.1.6	390-2	August 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012	
3.1.7	390-2	November 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012	
3.2	390-2	March 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.2.1	390-2	April 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	



Table 2-4. Intel® Manycore Platform Software Stack Revision Guide

Revision	Flash Revision	Customer Release Date	Supported Stepping	Supported Operating Systems	Workaround for Errata
3.2.3	390-2	May 2014	BO B1 CO	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.2.4	390-2	August 2014	BO B1 CO	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.2.5	390-2	November 2014	BO B1 CO	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.2.6	391-2	December 2014	BO B1 CO	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.2.7	391-2	March 2015	BO B1 CO	Red Hat Enterprise Linux* 6.0, 6.1, 6.2, 6.3, 6.4, 6.5 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.3	390-2	July 2014	BO B1 CO	Red Hat Enterprise Linux* 6.2, 6.3, 6.4, 6.5, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	




Table 2-4. Intel® Manycore Platform Software Stack Revision Guide

Revision	Flash Revision	Customer Release Date	Supported Stepping	Supported Operating Systems	Workaround for Errata
3.3.1	390-2	September 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.2, 6.3, 6.4, 6.5, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.3.2	390-2	October 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.2, 6.3, 6.4, 6.5, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.3.3	391-2	December 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.2, 6.3, 6.4, 6.5, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.3.4	391-2	March 2015	B0 B1 C0	Red Hat Enterprise Linux* 6.2, 6.3, 6.4, 6.5, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.4.0	390-2	September 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.3, 6.4, 6.5, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	



Table 2-4. Intel® Manycore Platform Software Stack Revision Guide

Revision	Flash Revision	Customer Release Date	Supported Stepping	Supported Operating Systems	Workaround for Errata
3.4.1	390-2	October 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.3, 6.4, 6.5, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.4.2	390-2	December 2014	B0 B1 C0	Red Hat Enterprise Linux* 6.3, 6.4, 6.5, 6.6, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	
3.4.3	391-2	February 2015	B0 B1 C0	Red Hat Enterprise Linux* 6.3, 6.4, 6.5, 6.6, 7.0 SUSE Linux Enterprise Server* 11 SP2, SP3 Microsoft Windows* 7 Enterprise SP1 Microsoft Windows* 8 Enterprise Microsoft Windows* 8.1 Enterprise Microsoft Windows* Server 2008 R2 SP1 Microsoft Windows* Server 2012 Microsoft Windows* Server 2012 R2	

Note: The latest Intel® Manycore Platform Software Stack can be found in the files available at <http://software.intel.com/en-us/articles/intel-manycore-platform-software-stack-mpss>. This site contains the firmware revision described in this table, along with the readme files.



3.0 Summary Table of Changes

The table included in this section indicates the errata, Specification Changes, Specification Clarifications, or Document Changes which apply to the Intel® Xeon Phi™ coprocessor. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding errata through documentation or specification changes as noted.

3.1 Affected Stepping Column:

X: This erratum applies to this stepping.

Blank: This erratum is fixed, or does not exist, in the listed stepping.

3.2 Status Column:

No Silicon Fix: Root caused to an erratum that will not be fixed in silicon. Note that a fix may be provided in an Intel® MPSS release. Please see the [Intel® Manycore Platform Software Stack Revision Guide](#) for more information.

Plan Fix: Root caused to an erratum and will be fixed in a future stepping.

Fixed: Root caused to an erratum and has been fixed in a subsequent stepping.

Spec Change: Root caused to a specification error that will be updated.

Investigating: A root cause has not been determined.

Third Party: Root caused to a board, software, driver, BIOS, or third party silicon issue.

3.3 Change Bar

Change bar on outside margins indicates this erratum is either new or modified from the previous version of the document.

Errata Summary Table (Sheet 1 of 3)

Number	Stepping			Status	Description
	B-0	B-1	C-0		
CD1.	X	X	X	No Silicon Fix	Not All Possible Instructions Prefixed With C4 And C5 Are Decoded
CD2.	X	X	X	No Silicon Fix	Writing to Reserved Bits in Certain Registers Does Not Result in a #GP
CD3.	X	X	X	No Silicon Fix	TLBs Will Not be Invalidated Upon a Page Fault
CD4.	X	X	X	No Silicon Fix	LZCNT Instruction With an Unsupported ModR/M Byte May Lead to an Incorrect Register Value
CD5.	X	X	X	No Silicon Fix	#NM Exception May be Raised on Certain EVEX Prefixed Scalar Instructions



Errata Summary Table (Sheet 2 of 3)

Number	Stepping			Status	Description
	B-0	B-1	C-0		
CD6.	X	X	X	No Silicon Fix	Spurious Stack Fault May Occur on Misaligned Stack Address
CD7.	X	X	X	No Silicon Fix	Bit Unexpectedly Set When DMA Channel is Disabled in Specific Way
CD8.	X	X	X	No Silicon Fix	PCIe Ordering Violation When a Small Posted Write Precedes an Interrupt Message
CD9.	X	X	X	No Silicon Fix	VPREFETCH And CLEVICT Instructions May Not Perform as Expected
CD10.	X	X	X	No Silicon Fix	Spurious MCA Error on Boot
CD11.	X	X	X	No Silicon Fix	PCIe Link Width May Degrade After Reset or PCIe Generation Transition
CD12.	X	X	X	No Silicon Fix	Simultaneous MCA Errors Will be Reported When MCA Reporting is Disabled
CD13.	X	X	X	No Silicon Fix	Accessing Illegal MMIO Addresses May Cause Hang
CD14.	X	X	X	No Silicon Fix	Pairing of Either JKZD or JKNZD Instructions May Fail to Raise Expected #NM Exception
CD15.	X	X	X	No Silicon Fix	Unexpected #GP Fault on a #UD Fault Instruction
CD16.	X	X	X	No Silicon Fix	Incorrect Ordering When DMA Status Descriptor Precedes an Interrupt
CD17.	X	X	X	No Silicon Fix	Unexpected DMA Ordering Between Two Status Descriptor Posted Writes
CD18.	X	X	X	No Silicon Fix	Unexpected #UD Fault on Multi-byte NOP Instruction
CD19.	X	X	X	No Silicon Fix	MXCSR Register May Contain Unexpected Values Following a Faulting Vector Instruction
CD20.	X	X	X	No Silicon Fix	L1 Entry Earlier Than Expected
CD21.	X	X	X	No Silicon Fix	Double Fault May Not be Signaled
CD22.	X	X		Fixed	Spurious Errors When Not Operating in Common Clock Mode
CD23.	X	X		Fixed	Incorrect Value Reported in IOAPIC Maximum Redirection Field
CD24.	X	X	X	No Silicon Fix	Erroneously Detect 16 Lanes After First Detect.Active
CD25.	X	X	X	No Silicon Fix	Specification Violation For Rx Return Loss in Common Mode
CD26.	X	X	X	No Silicon Fix	VPACKSTOREH Instruction Paired With Memory Operation
CD27.	X	X		Fixed	Hang When Package C6 State is Enabled
CD28.	X	X	X	No Silicon Fix	FNCLEX Does Not Clear Unmasked Exception
CD29.	X	X	X	No Silicon Fix	Livelock on Branch to HLT or Branch Before HLT
CD30.	X	X	X	No Silicon Fix	Data Breakpoint May be Taken on Initiating Instruction
CD31.	X	X	X	No Silicon Fix	Coprocessor Stuck in Lower P-State
CD32.	X	X	X	Fixed	NMI During DPC3 Transition
CD33.	X	X	X	No Silicon Fix	TCP/IP NFS Low Performance
CD34.	X	X	X	No Silicon Fix	Application Threads Migrate With Coprocessor OS Scheduler
CD35.	X	X	X	No Silicon Fix	Host CPU P-state Transitions Affect PCIe Transfer Rate
CD36.	X	X	X	No Silicon Fix	Spurious Bad TLP or Bad DLLP Correctable Error Status on Electrical Idle
CD37.	X	X	X	No Silicon Fix	GBOX MMIO Write Failure
CD38.	X	X		Fixed	SMC Return Erroneous Data and SMC Firmware Update Failure
CD39.	X	X	X	No Silicon Fix	Disabling Active DMA Channel or Re-enabling DMA Channel Without Setting DRARLO or DRARHI Leads to Unexpected System Behavior



Errata Summary Table (Sheet 3 of 3)

Number	Stepping			Status	Description
	B-0	B-1	C-0		
CD40.	X	X		Fixed	DBOX MMIO Read/Write Failure During Frequency Transitions
CD41.	X	X	X	No Silicon Fix	Data Breakpoint Triggered by a Thread May Signal Breakpoint in Other Threads
CD42.	X	X	X	No Silicon Fix	Frequency Ratio Value in CURRENTRATIO Register Incorrect Immediately After Frequency Change
CD43.	X	X	X	No Silicon Fix	SMC SMBus Auto-negotiation Does Not Function on Shared Bus
CD44.	X	X	X	No Silicon Fix	MEMCPY Descriptor May be Executed When Destination Address is in SBOX or DBOX MMIO
CD45.	X	X	X	No Silicon Fix	DMA Descriptor Fetch Error May Flag Source Address Error
CD46.	X	X	X	No Silicon Fix	SBOX System Interrupt Register Read Timeout
CD47.	X	X	X	No Silicon Fix	Spurious PCIe Interrupt While Coprocessor is in Deep-pC3 or pC6
CD48.	X	X	X	No Silicon Fix	Invalid Link Training Sequence After L1 Entry Request
CD49.	X	X	X	No Silicon Fix	DMA Engine Does Not Flag DESCR_ADDR_ERR When DRAR is Not Cacheline Aligned
CD50.	X	X	X	No Silicon Fix	Memory Reads or Writes Above Top of Memory Not Invalidated
CD51.	X	X	X	No Silicon Fix	Host Targeted Interrupt May Prevent Illegal Access MCA Event Signal
CD52.	X	X	X	No Silicon Fix	DMA Descriptors May Execute Out of Order if DMA Status Descriptor is Used
CD53.	X	X	X	No Silicon Fix	SBOX I/O APIC Does Not Handle Low Priority Interrupts
CD54.	X	X	X	No Silicon Fix	Incorrect Sensor Data From Coprocessor SMC
CD55.	X	X	X	No Silicon Fix	PCIe Read Bandwidth Reduction Behind PCIe Switch
CD56.	X	X	X	No Silicon Fix	Coprocessor Hang During PC3 or DPC3 Exit

Specification Changes


Number	SPECIFICATION CHANGES
CD1.	System Software Developer Guide: Appendix: SBOX Control Register List:

Specification Clarifications

Number	SPECIFICATION CLARIFICATIONS
CD1.	Clarification of Support For Receiving an Unequal Number of SKP Symbols Across PCI Express Lanes

Documentation Changes

Number	DOCUMENTATION CHANGES
	There are no Documentation Changes at this time.



4.0 Errata

CD1. Not All Possible Instructions Prefixed With C4 And C5 Are Decoded

Problem: The coprocessor fails to decode all possible C4/C5-prefixed instructions. The shorter 8-bit displacement form of JKZD/JKNZD may not be used with the C5-prefix. For "mask k, k" instructions, the coprocessor will only allow the C5-prefix to be used.

Implication: In the case of JKZD/JKNZD, only the one-byte longer C4-prefix can be used; this results in a larger code size. For 'mask k, k' instructions, the longer C4-prefix will not be recognized.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD2. Writing to Reserved Bits in Certain Registers Does Not Result in a #GP

Problem: Writing to reserved bits 63:40 in either IA32_PerfCntr0 (20h) or IA32_PerfCntr1 (21h), and reserved bits 31:4 in SBOX_EMON_CNT_OVFL will not generate a #GP (general protection fault) as expected.

Implication: Due to this erratum, any value written to reserved bits are ignored and will not generate a #GP.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD3. TLBs Will Not be Invalidated Upon a Page Fault

Problem: The TLBs (Translation Lookaside Buffer) will not be invalidated upon page fault. This is not consistent with other recent Intel(R) architecture products, in which the hardware automatically invalidates the TLBs.

Implication: If the coprocessor operating system page fault handler does not perform a TLB invalidation, the page fault may recur.

Workaround: Software must invalidate the TLB's entry upon page fault.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).


CD4. LZCNT Instruction With an Unsupported ModR/M Byte May Lead to an Incorrect Register Value

Problem: If the LZCNT instruction contains an unsupported ModR/M byte, its execution will trigger #UD fault handler. However, due to this erratum, the instruction following the LZCNT may be executed prior to processing the #UD exception.

Implication: The instruction following LZCNT may be executed, which may lead to an incorrect register value.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

**CD5. #NM Exception May be Raised on Certain EVEX Prefixed Scalar Instructions**

Problem: The EVEX prefixed versions of the CLEVICT and VPREFETCH instructions will generate an incorrect #NM (Device Not Available) exception when the CRO.TS bit is set (bit 3).

Implication: Due to this erratum, an unexpected exception may be generated.

Workaround: None identified. This erratum can be avoided by either implementing an #NM fault handler in the coprocessor operating system, not setting the CRO.TS bit, or not using the EVEX prefixed versions of the CLEVICT and VPREFETCH instructions.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD6. Spurious Stack Fault May Occur on Misaligned Stack Address

Problem: In instances where the coprocessor operating system does not automatically align the stack address and the stack address is misaligned so that its address crosses the 4GB boundary, a spurious stack fault may occur.

Implication: An unexpected stack fault may be generated.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD7. Bit Unexpectedly Set When DMA Channel is Disabled in Specific Way

Problem: Disabling an active DMA (Direct Memory Access) channel by clearing its enable bit in the DCR register (0xA280) causes the DSTAT.Q bit (bit 28 of the respective DMA_DSTAT register) to be unexpectedly set.

Implication: The DSTAT.Q bit indicates that the DMA channel is quiesced, which is not correct. Applications that look to the DSTAT.Q bit to determine DMA channel quiesce status will be affected by this erroneous value.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD8. PCIe Ordering Violation When a Small Posted Write Precedes an Interrupt Message

Problem: When any posted write with a payload that is less than 16 bytes precedes an Interrupt Message, ordering between the interrupt and the posted write may be incorrect.

Implication: Features such as DMA transfers and peer-to-peer may not work correctly if an Interrupt Message is generated immediately after a posted write with a payload of less than 16 bytes.

Workaround: There are two situations where a workaround is needed to circumvent this erratum:

1) When a coprocessor hardware thread is sending a message to the host CPU and the message consists of two parts, part one being data written to system memory and part two being a doorbell write that sends an interrupt configured as an INTx (as opposed to MSI). In this situation, the coprocessor hardware thread will need to insert a write to system memory between the two parts.

2) When using DMA with system memory as the destination and a status descriptor is used to generate the interrupt to the host, a second status descriptor with the interrupt bit turned off must be inserted before the status descriptor that generates the interrupt.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD9. VPREFETCH And CLEVICT Instructions May Not Perform as Expected

Problem: If segment type is configured as "expand-down" in 32 bit mode, the VPREFETCH and CLEVICT instructions using either the GS or FS segment prefixes may be unexpectedly dropped when in 64 bit mode.



Implication: A noticeable reduction in performance may be observed.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD10. Spurious MCA Error on Boot

Problem: A coprocessor with 12 GDDR5 channels may trigger a spurious MCA (Machine Check Architecture) error (GBOX_MBOX_MCA_STATUS_LO) when it boots.

Implication: An unexpected MCA error may be logged. There are no functional implications.

Workaround: None identified. This spurious MCA error can be safely cleared on boot.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD11. PCIe Link Width May Degrade After Reset or PCIe Generation Transition

Problem: The PCIe link width may degrade after a warm reset, hot reset, or when transitioning between PCIe generations.

Implication: The link width will be reset to a lower width than expected. For example, it might be set to 4 lanes when it was initially set to 16 lanes.

Workaround: None identified.

Status: Fix planned in firmware update. Please refer to the [Firmware Revision](#) table.

CD12. Simultaneous MCA Errors Will be Reported When MCA Reporting is Disabled

Problem: When simultaneous MCA (Machine Check Architecture) errors occur, the errors will be reported and associated interrupts signaled if MCA reporting is disabled via the MCX_CTL_LO register.

Implication: Unexpected MCA errors will be reported when the reporting of MCA errors is disabled.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD13. Accessing Illegal MMIO Addresses May Cause Hang

Problem: Accessing illegal MMIO addresses (in the range of 0x800710000-0x80078FFFF) may not return 0 as expected.

Implication: The coprocessor may hang.

Workaround: Software using 2 MB pages may inadvertently access these illegal MMIO addresses if a valid page translation exists. Therefore using 4 KB or 64 KB pages to map the valid MMIO space is recommended. If the coprocessor hangs as a result of accessing these addresses, a reset of the host will be required in order to recover from the hang.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).


CD14. Pairing of Either JKZD or JKNZD Instructions May Fail to Raise Expected #NM Exception

Problem: If the JKZD or JKNZD instructions are paired with any other instruction and CR0.TS is set, they may fail to raise an expected #NM (Device Not Available) exception.

Implication: The coprocessor operating system will not be able to perform lazy context save and restore of mask registers. Vector data registers are not affected.

Workaround: The coprocessor operating system should explicitly save and restore mask registers on any context switch.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

**CD15. Unexpected #GP Fault on a #UD Fault Instruction**

Problem: When there is a #UD faulting instruction on the border of a canonical address space when the instruction length is 15 bytes, a spurious #GP fault may be generated.

Implication: The #GP fault will be given a higher priority than the #UD fault.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD16. Incorrect Ordering When DMA Status Descriptor Precedes an Interrupt

Problem: Executing a DMA status descriptor with an interrupt enabled may result in incorrect ordering between the status descriptor write and associated Interrupt Marker Message.

Implication: The Interrupt Marker Message may arrive before the data associated with the DMA status descriptor.

Workaround: To work around this erratum, use two status descriptors (one after the other) to the same location with the Interrupt bit enabled on the second status descriptor.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD17. Unexpected DMA Ordering Between Two Status Descriptor Posted Writes

Problem: When two consecutive DMA status descriptors are executed, and the first status descriptor has the interrupt bit set, the second status descriptor's write may arrive before the interrupt for the first status descriptor.

Implication: This will result in unexpected ordering between the status descriptors and the interrupt associated with the first status descriptor.

Workaround: Software should expect that the second status descriptor write may arrive before the interrupt associated with the first status descriptor write.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD18. Unexpected #UD Fault on Multi-byte NOP Instruction

Problem: The coprocessor expects the multi-byte NOP to have the MODRM.reg field set to 0. Previous x86 programmer's reference manuals may indicate that setting the MODRM.reg field is not necessary for a multi-byte NOP instruction, such as "0F 01".

Implication: Software which implements a multi-byte NOP while failing to set the MODRM.reg field to 0 will generate an unexpected #UD fault.

Workaround: This erratum can be worked around by implementing a multi-byte NOP as follows: "0F 1F /0". This will effectively set the MODRM.reg field to 0.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD19. MXCSR Register May Contain Unexpected Values Following a Faulting Vector Instruction

Problem: If a vector instruction takes a fault while fetching an operand or memory, the instruction will not complete and will erroneously update the values in the MXCSR register.

Implication: The MXCSR register will contain unexpected values when examined by a fault handler.

Workaround: Fault handlers should expect the MXCSR register to contain unexpected values when examined after a faulting vector instruction.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

**CD20. L1 Entry Earlier Than Expected**

Problem: The coprocessor may request PCIe L1 entry 2 microseconds earlier than what is specified in the L1_ENTRY_TIMER register (0x4124).

Implication: This will result in an L1 entry sooner than expected. There are no functional implications.

Workaround: Software expecting a L1 entry request according to the value set in L1_ENTRY_TIMER should expect the request to come 2 microseconds sooner.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD21. Double Fault May Not be Signaled

Problem: If the coprocessor takes a second fault while handling a first fault, the coprocessor may ignore the first fault and signal the second fault's handler.

Implication: The second fault is taken, and a double fault is not signaled.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD22. Spurious Errors When Not Operating in Common Clock Mode

Problem: When the coprocessor is not operating in common clock mode, the coprocessor may incorrectly signal Bad DLLP or Bad TLP errors.

Implication: The PCIe link will transition in to and out of recovery automatically and signal these spurious errors, which may cause management software to erroneously indicate that the coprocessor is faulty.

Workaround: None identified

Status: Fixed. For the steppings affected, see the [Errata Summary Table](#).

CD23. Incorrect Value Reported in IOAPIC Maximum Redirection Field

Problem: The IOAPIC Maximum Redirection Field (23:16) on the coprocessor incorrectly reports a value of 24.

Implication: Software that depends on this field to determine the number of DMA channels can not assign an IRQ to the RDMSR7 (Remote DMA Register 7, 0xB19C) SBOX register. This will prevent RDMA from functioning to an eighth coprocessor.

Workaround: None identified.

Status: Fixed. For the steppings affected, see the [Errata Summary Table](#).

CD24. Erroneously Detect 16 Lanes After First Detect.Active

Problem: During the training sequence, if the coprocessor does not detect 16 lanes during the first Detect.Active, the coprocessor will erroneously detect 16 lanes regardless of whether the PCIe connector has 16 lanes available.

Implication: The coprocessor may enter Polling.Active before RX Termination is enabled on the link partner. If the PCIe connector has less than 16 lanes, the coprocessor may send training sequences on 16 lanes. The link will eventually train correctly.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD25. Specification Violation For Rx Return Loss in Common Mode

Problem: The coprocessor Rx return loss for common mode may be higher than the PCI Express Gen2 specification for certain frequency ranges.

Implication: There is a violation of the PCI Express Gen2 Base Specification for Rx return loss. No functional implications have been observed.



Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD26. VPACKSTOREH Instruction Paired With Memory Operation

Problem: When a VPACKSTOREH instruction is used with an uncacheable memory type, has a 64-byte aligned address, and is paired with another memory operation, the paired memory operation may not properly execute.

Implication: Unpredictable results may occur due to this erratum.

Workaround: If there is a VPACKSTOREL instruction associated with the VPACKSTOREH instruction (to effect storing unaligned data), then order them such that the VPACKSTOREL instruction is immediately, lexically after the VPACKSTOREH instruction. Otherwise, assure some other non-memory-accessing instruction is immediately lexically after the VPACKSTOREH instruction, such as a NOP.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD27. Hang When Package C6 State is Enabled

Problem: The coprocessor may hang when it is allowed to transition to the Package C6 state.

Implication: The coprocessor is not able to take advantage of the lowest power state. Deep C3 state and memory power states are not affected.

Workaround: Do not enable Package C6 transition for the coprocessor.

Status: Fixed. For the steppings affected, see the [Errata Summary Table](#).

CD28. FNCLEX Does Not Clear Unmasked Exception

Problem: When the FNCLEX instruction is executed after an unmasked floating-point exception has occurred, it will not properly clear the unmasked exception as expected.

Implication: If the next arithmetic floating point instruction following the FNCLEX instruction raises a new exception, it may take the wrong exception or get the wrong result.

Workaround: Instead of FNCLEX, software should execute FNSAVE, clear the bits in the memory image that map to FSW, and execute FRSTOR.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD29. Livelock on Branch to HLT or Branch Before HLT

Problem: A branch to the HLT (halt) instruction or a branch before the HLT instruction may cause a livelock condition.

Implication: Due to this erratum, the coprocessor may livelock until another thread on the same core takes an interrupt.

Workaround: Software should not branch to the HLT instruction or have the HLT instruction following a branch.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD30. Data Breakpoint May be Taken on Initiating Instruction

Problem: Data breakpoints may be taken on the instruction that initiated the break point instead of the following instruction as expected.

Implication: Due to this erratum, the break point will happen one instruction earlier than expected.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD31. Coprocessor Stuck in Lower P-State

Problem: The coprocessor may become stuck in the lowest p-state following a thermal throttling event.



Implication: The coprocessor will not return to the highest p-state, and performance may be impacted.

Workaround: None identified.

Status: Fixed in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#).

CD32. NMI During DPC3 Transition

Problem: Transitions to and from DPC3 (Deep Package C3) may cause NMI (Non-Maskable Interrupts).

Implication: Platforms may behave differently on account of the NMI.

Workaround: DPC3 transitions have been disabled by default in the coprocessor OS.

Status: Fixed in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#). For the steppings affected, see the [Errata Summary Table](#).

CD33. TCP/IP NFS Low Performance

Problem: Performance of NFS over the virtual ethernet interface with TCP may be lower than expected.

Implication: Applications which depend on high NFS performance while NFS is mounted over TCP may be impacted.

Workaround: Using UDP with NFS may result in better performance.

Status: Fix Planned in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#). For the steppings affected, see the [Errata Summary Table](#).

CD34. Application Threads Migrate With Coprocessor OS Scheduler

Problem: The scheduler in the coprocessor OS may migrate application threads between cores when a specific affinity is not specified.

Implication: This can lead to variable performance between subsequent runs of the same application

Workaround: This can be worked around by configuring application thread affinity.

Status: Fix Planned in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#). For the steppings affected, see the [Errata Summary Table](#).

CD35. Host CPU P-state Transitions Affect PCIe Transfer Rate

Problem: Host CPU p-state transitions to a lower frequency p-state can lead to a reduction in PCIe transfer bandwidth.

Implication: Applications may exhibit irregular PCIe transfer rates.

Workaround: This can be worked around by disabling host CPU p-state transitions.


Status: Fixed in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#)

CD36. Spurious Bad TLP or Bad DLLP Correctable Error Status on Electrical Idle

Problem: If L0s is enabled, a spurious Bad TLP or Bad DLLP correctable error may be signaled after several transitions to Electrical Idle on the Rx lanes.

Implication: The coprocessor may erroneously indicate a link integrity problem. Other transitions to Electrical Idle are under investigation.

Workaround: To avoid these spurious correctable errors, disable ASPM L0s.



Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD37. GBOX MMIO Write Failure

Problem: A write to a GBOX MMIO register may fail when there is heavy memory write traffic.

Implication: Due to this erratum, EMON, PM (Power Management), and MCA status register writes to the GBOX may not occur as expected.

Workaround: Software can write to the GBOX MMIO register then read it to verify, and perform another write if it was not successful. In the case of write-only registers, writing the register multiple times will reduce the chance of this erratum occurring.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD38. SMC Return Erroneous Data and SMC Firmware Update Failure

Problem: Communication between the coprocessor and the SMC (System Manageability Controller) may occasionally fail. For information on identifying coprocessors that may be affected by this erratum, see the [Intel® Xeon Phi™ Coprocessor Identification](#) table.

Implication: Due to this erratum, applications or platform BMC code which depend on data returned by the SMC may receive erroneous data (typically all zeroes) and SMC firmware updates may not complete successfully, leaving the SMC in an unusable state until a successful update is achieved.

Workaround: None identified.

Status: Fixed. For the steppings affected, see the [Errata Summary Table](#).

CD39. Disabling Active DMA Channel or Re-enabling DMA Channel Without Setting DRARLO or DRARHI Leads to Unexpected System Behavior

Problem: The coprocessor DMA engine does not support disabling a DMA channel while that channel is active, or re-enabling the DMA channel without writing to channel's respective DRARLO or DRARHI register.

Implication: Due to this erratum, when an active DMA channel is disabled or a channel is re-enabled without writing to channel's respective DRARLO or DRARHI register, unexpected system behavior may occur.

Workaround: None identified. A DMA channel should only be disabled after the active bit in the channel's respective DSTAT register (bit 29) been cleared. Software should set the channel's DRARLO or DRARHI register before re-enabling the channel.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD40. DBOX MMIO Read/Write Failure During Frequency Transitions

Problem: In certain rare circumstances, coprocessor frequency transitions may interfere with DBOX MMIO register accesses.

Implication: Due to this erratum, a read or write operation on a DBOX MMIO register may fail.


Workaround: Performing a read operation on a MCLK register in DBOX (for example, LPCECNT0 (0xE01C)) prior to performing a read or write operation on a DBOX MMIO register may reduce the chance of this occurring.

Status: Fixed. For the steppings affected, see the [Errata Summary Table](#).

CD41. Data Breakpoint Triggered by a Thread May Signal Breakpoint in Other Threads

Problem: When using data breakpoints to debug multithreaded software and various breakpoints have been set across all threads, a breakpoint encountered by one thread may also be signaled by other threads in certain rare circumstances.

Implication: Due to this erratum, a data breakpoint enabled for a particular thread will still be triggered as expected by that thread, but may also be signaled in threads where the breakpoint was not hit.



Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD42. Frequency Ratio Value in CURRENTRATIO Register Incorrect Immediately After Frequency Change

Problem: After a coprocessor frequency change, there is a short delay until the new frequency ratio is reflected in the CURRENTRATIO register (MMIO offset 0x402C).

Implication: Software that reads from the CURRENTRATIO register immediately after a frequency change has occurred may not read a value that represents correct coprocessor frequency ratio. Frequency changes can occur from within the coprocessor OS kernel, or from an external throttle event originating from the SMC (System Management Controller) or platform BMC.

Workaround: None identified. In order to greatly reduce the likelihood of encountering this erratum, it is recommended that software perform multiple reads of the CURRENTRATIO register, stopping when the value returned has been consistent for two or more subsequent reads.

Status: Fixed in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#)

CD43. SMC SMBus Auto-negotiation Does Not Function on Shared Bus

Problem: SMC (System Management Controller) SMBus address auto-negotiation does not function as expected when there are multiple coprocessors sharing the same bus.

Implication: Due to this erratum, certain coprocessors may be unresponsive when accessed out-of-band over SMBus. This erratum does not affect coprocessors that are installed behind a SMBus switch (mux).

Workaround: None identified.

Status: Fixed in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#)

CD44. MEMCPY Descriptor May be Executed When Destination Address is in SBOX or DBOX MMIO

Problem: Error checking logic may not correctly flag an error (DEST_ADDR_ERR) for a MEMCPY descriptor that has a destination address in the SBOX or DBOX MMIO region.

Implication: Due to this erratum, the descriptor is executed, which may result in undefined behavior.

Workaround: None identified. It is recommended that software should not program a MEMCPY descriptor with a destination address in the SBOX or DBOX MMIO region.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#).

CD45. DMA Descriptor Fetch Error May Flag Source Address Error

Problem: Error signaling logic may incorrectly flag a DMA descriptor fetch error to host memory as a DMA data fetch error.


Implication: Due to this erratum, if the DMA descriptor fetch fails, a source address error (SRC_ADDR_ERR) is flagged instead of a descriptor address error (DESCR_ADDR_ERR).

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#)

CD46. SBOX System Interrupt Register Read Timeout

Problem: Under rare circumstances, reads to System Interrupt registers in the coprocessor SBOX may timeout.



Implication: Due to this erratum, debug software expecting to utilize System Interrupt registers in the coprocessor SBOX may be impacted.

Workaround: None identified.

Status: Fix Planned in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#). For the steppings affected, see the [Errata Summary Table](#).

CD47. Spurious PCIe Interrupt While Coprocessor is in Deep-pC3 or pC6

Problem: If the coprocessor is in the Deep-pC3 or pC6 sleep state and receives a malformed TLP, a spurious interrupt may be sent in addition to the Completer Abort (CA) Status.

Implication: When this occurs, the coprocessor will request the host driver to take actions to transition the coprocessor back to full power state.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#)

CD48. Invalid Link Training Sequence After L1 Entry Request

Problem: The coprocessor may send invalid Link Training Sequences when exiting up from a short L1 transition.

Implication: Due to this erratum, the PCIe root complex may receive invalid Link Training Sequences from the coprocessor when exiting from a short L1 transition. The coprocessor will start sending valid Link Training Sequences once the root complex begins to send Link Training Sequences.

Workaround: None identified.

Status: Fix Planned in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#). For the steppings affected, see the [Errata Summary Table](#).

CD49. DMA Engine Does Not Flag DESC_ADDR_ERR When DRAR is Not Cacheline Aligned

Problem: The DMA engine descriptor ring will not flag a descriptor address error (DESCR_ADDR_ERR) when the lower 6 bits of DRAR (Descriptor Ring Attributes Register) are not cacheline aligned.

Implication: Due to this erratum, failure to align the lower 6 bits of DRAR to a cacheline may result in unexpected behavior.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#)

CD50. Memory Reads or Writes Above Top of Memory Not Invalidated

Problem: The coprocessor may not invalidate memory reads or writes to an area outside of physical memory. Zeros are normally returned for a read to an invalid memory location, while a write normally results in the write being dropped.

Implication: If multiple reads or writes to an area outside of physical memory occur, the coprocessor may hang.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#)

CD51. Host Targeted Interrupt May Prevent Illegal Access MCA Event Signal

Problem: An Illegal Access MCA event (0x0020h) occurring at the same time as an interrupt targeted to the host may prevent the Illegal Access MCA event from being signaled.



Implication: Due to this erratum, an Illegal Access MCA event may not be signaled when the illegal access occurs.

Workaround: None identified.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#)

CD52. DMA Descriptors May Execute Out of Order if DMA Status Descriptor is Used

Problem: A DMA descriptor following a DMA status descriptor that is targeting SBOX register space will not wait for the DMA status descriptor to complete.

Implication: Due to this erratum, if software uses a status descriptor targeting SBOX registers, the following DMA descriptors will start executing before the status descriptor completes.

Workaround: None identified. It is recommended that software uses two general purpose descriptors instead of a single status descriptor when targeting SBOX register space with a 64 bit write.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#)

CD53. SBOX I/O APIC Does Not Handle Low Priority Interrupts

Problem: Low priority interrupts sent in broadcast or targeting the SBOX are not correctly handled by the SBOX I/O APIC.

Implication: When the SBOX I/O APIC receives a low priority interrupt, unexpected behavior may occur.

Workaround: None identified. It is recommended that low priority interrupts target a CPU or address a set of local APICs (excluding SBOX) defined as a logical group using logical destination mode.

Status: No silicon fix planned. For the steppings affected, see the [Errata Summary Table](#)

CD54. Incorrect Sensor Data From Coprocessor SMC

Problem: The coprocessor SMC may return incorrect sensor data when queried by the platform BMC, despite the sensor on the coprocessor being present and functional. The coprocessor SMC will appropriately indicate that this sensor data is incorrect by setting the Reading/State Unavailable flag as part of the IPMI Get Sensor Reading Command response to the platform BMC.

Implication: Due to this erratum, the platform BMC may take action based on this incorrect sensor data if the platform BMC does not properly handle the Reading/State Unavailable flag. For example, the coprocessor SMC may return a value of '0' for coprocessor die temperature, which may cause the platform BMC to incorrectly reduce fan speeds.

Thermal management features on the coprocessor itself, such as throttling and THERMTRIP#, are not affected by this erratum.

Workaround: None identified. It is recommended that the platform BMC detect and properly handle the Reading/State Unavailable flag of the IPMI Get Sensor Reading Command.


Status: Fixed in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#). For the steppings affected, see the [Errata Summary Table](#).

CD55. PCIe Read Bandwidth Reduction Behind PCIe Switch

Problem: PCIe read requests originating from the coprocessor may experience a degradation in bandwidth when the coprocessor is connected to a PCIe switch.

Implication: When this occurs, reduced PCIe read bandwidth may be observed when operating the coprocessor behind a PCIe switch.

Workaround: None identified.



Status: No fix planned. Workaround implemented in Intel(R) MPSS flash to improve bandwidth. Please refer to the [Intel® Manycore Platform Software Stack Revision Guide](#) for a list of Intel(R) MPSS releases that include the workaround.


CD56. Coprocessor Hang During PC3 or DPC3 Exit

Problem: Coprocessors exiting the PC3 (Package C3) or DPC3 (Deep Package C3) power states may hang. While this behavior may occur on any exit from PC3 or DPC3, it has been observed to occur after anywhere from a few hours to a few days (on average, millions of transitions). It is not associated with any specific SKU.

Implication: When the hang occurs due to this erratum, users will not be able to connect to or execute applications on the coprocessor until the host platform has been reset.

Workaround: This erratum can be avoided by disabling the PC3 and DPC3 power management states. Please refer to the Intel(R) MPSS readme documentation for instructions on disabling these power management states. For coprocessors that support the PC6 power state, disabling DPC3 will also disable the PC6 functionality.

Status: Fixed in Intel® MPSS. For a list of Intel® MPSS releases affected, see the [Intel® Manycore Platform Software Stack Revision Guide](#). For the steppings affected, see the [Errata Summary Table](#).



5.0 Specification Changes

The Specification Changes listed in this section apply to the following documents:

- *Intel® Xeon Phi™ Coprocessor Datasheet*
- *Intel® Xeon Phi™ Coprocessor System Software Developers Guide*
- *Intel® Xeon Phi™ Coprocessor Instruction Set Architecture Reference Manual*
- *PCIe® Base 2.1 Specification*
- *Intel® Manycore Platform Software Stack Readme*

CD1. **System Software Developer Guide: Appendix: SBOX Control Register List:**

The Reset Domain for the DCHERRMSK registers is “Hot” instead of “Sticky”, indicating that these DMA Channel Error Mask registers are reset after a PCIe hot reset.



6.0 Specification Clarifications

The Specification Clarifications listed in this section apply to the following documents:

- *Intel® Xeon Phi™ Coprocessor Datasheet*
- *Intel® Xeon Phi™ Coprocessor System Software Developers Guide*
- *Intel® Xeon Phi™ Coprocessor Instruction Set Architecture Reference Manual*
- *PCIe® Base 2.1 Specification*
- *Intel® Manycore Platform Software Stack Readme*

CD1. Clarification of Support For Receiving an Unequal Number of SKP Symbols Across PCI Express Lanes

The Intel® Xeon Phi™ coprocessor does not support Repeaters on its Link that send Skip Ordered-Sets (SOS) with an unequal number of SKP Symbols across two or more Lanes. In response to receiving such a SOS on the Link, the following behavior may be seen:

- 1) Bad DLLP or Bad TLP Correctable errors may be reported by the coprocessor.
- 2) The Root Port may initiate a Link Layer Replay.
- 3) Subsequent TLP's may not be correctly interpreted by the coprocessor, resulting in a Replay Timer Timeout Correctable error reported by the Root Port.
- 4) The Root Port may continue to Replay and report a REPLAY_NUM Rollover Correctable error and transition the Link into Recovery.
- 5) The coprocessor receiver logic will then recover and be able to correctly receive packets from the Link.



7.0 Documentation Changes

The Documentation Changes listed in this section apply to the following documents:

- *Intel® Xeon Phi™ Coprocessor Datasheet*
- *Intel® Xeon Phi™ Coprocessor System Software Developers Guide*
- *Intel® Xeon Phi™ Coprocessor Instruction Set Architecture Reference Manual*
- *PCIe® Base 2.1 Specification*
- *Intel® Manycore Platform Software Stack Readme*

CD1. There are no Documentation Changes at this time.