



AMBA CHI Issue E.b Errata

Architecture & Technology Group

Document number: ARM AES 0061
Release Number: 4.0
Confidentiality: Non-Confidential
Date of Issue: 20-Sep-2022

© Copyright Arm Limited 2022. All rights reserved.

Abstract

This document includes clarifications and corrections to the CHI Issue E.b specification.

Contents

About this document	3
Release Information	3
References	3
1	Introduction
1.1	Classification of the change
2	New or Updated Errata
2.1	C597: Cancellation of CopyBack requests following overlapping snoop
2.2	D638: ReadReceipt not optional in certain transaction flows
2.3	C624: DVM payload encoding for Instruction cache invalidations
3	Previous Errata
3.1	C605: DBID values in Comp and DBIDResp messages that originate from difference sources in DWT flow have no relationship
3.2	C606: Data_Check and Check_Type property descriptions
3.3	D607: Typographical error in Allocating Read figure
3.4	D609: Typographical error in RN to HN Write request attribute values table
3.5	C612: Requirements for completer read response when MTE is not supported
3.6	C621: Typographical error in example WriteUniqueStash with Data Pull transaction flow
3.7	C628: Requirements for ReadOnceMakeInvalid when invalidating a Dirty copy

About this document

Release Information

The change history table lists the changes that have been made to this document.

Date	Version	Confidentiality	Change
03-Feb-2022	1.0	Confidential	First limited release
25-Apr-2022	2.0	Non-Confidential	First public release
25-Aug-2022	3.0	Confidential	Second limited release
20-Sep-2022	4.0	Non-Confidential	Second public release

References

This document refers to the following documents.

Ref	Document Number	Title
1	ARM IHI 0050E.b	AMBA® 5 CHI Architecture Specification

Proprietary Notice

Proprietary Notice

This document is **NON-CONFIDENTIAL** and any use by you is subject to the terms of this notice and the Arm AMBA Specification Licence set about below.

This document is protected by copyright and other related rights and the practice or implementation of the information contained in this document may be protected by one or more patents or pending patent applications. No part of this document may be reproduced in any form by any means without the express prior written permission of Arm. **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.**

Your access to the information in this document is conditional upon your acceptance that you will not use or permit others to use the information for the purposes of determining whether implementations infringe any third party patents.

THIS DOCUMENT IS PROVIDED "AS IS". ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, patents, copyrights, trade secrets, or other rights.

This document may include technical inaccuracies or typographical errors.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

This document consists solely of commercial items. You shall be responsible for ensuring that any use, duplication or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word "partner" in reference to Arm's customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

This document may be translated into other languages for convenience, and you agree that if there is any conflict between the English version of this document and any translation, the terms of the English version of the Agreement shall prevail.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners. Please follow Arm's trademark usage guidelines at <http://www.arm.com/company/policies/trademarks>.

Copyright © [2022] Arm Limited (or its affiliates). All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

LES-PRE-21451 version 2.2

AMBA SPECIFICATION LICENCE

THIS END USER LICENCE AGREEMENT ("LICENCE") IS A LEGAL AGREEMENT BETWEEN YOU (EITHER A SINGLE INDIVIDUAL, OR SINGLE LEGAL ENTITY) AND ARM LIMITED ("ARM") FOR THE USE OF ARM'S INTELLECTUAL PROPERTY (INCLUDING, WITHOUT LIMITATION, ANY COPYRIGHT) IN THE RELEVANT AMBA SPECIFICATION ACCOMPANYING THIS LICENCE. ARM LICENSES THE RELEVANT AMBA SPECIFICATION TO YOU ON CONDITION THAT YOU ACCEPT ALL OF THE TERMS IN THIS LICENCE. BY CLICKING "I AGREE" OR OTHERWISE USING OR COPYING THE RELEVANT AMBA SPECIFICATION YOU INDICATE THAT YOU AGREE TO BE BOUND BY ALL THE TERMS OF THIS LICENCE.

"LICENSEE" means You and your Subsidiaries.

"Subsidiary" means, if You are a single entity, any company the majority of whose voting shares is now or hereafter owned or controlled, directly or indirectly, by You. A company shall be a Subsidiary only for the period during which such control exists.

1. Subject to the provisions of Clauses 2, 3 and 4, Arm hereby grants to LICENSEE a perpetual, non-exclusive, non-transferable, royalty free, worldwide licence to:

(i) use and copy the relevant AMBA Specification for the purpose of developing and having developed products that comply with the relevant AMBA Specification;

(ii) manufacture and have manufactured products which either: (a) have been created by or for LICENSEE under the licence granted in Clause 1(i); or (b) incorporate a product(s) which has been created by a third party(s) under a licence granted by Arm in Clause 1(i) of such third party's AMBA Specification Licence; and

(iii) offer to sell, sell, supply or otherwise distribute products which have either been (a) created by or for LICENSEE under the licence granted in Clause 1(i); or (b) manufactured by or for LICENSEE under the licence granted in Clause 1(ii).

2. LICENSEE hereby agrees that the licence granted in Clause 1 is subject to the following restrictions:

(i) where a product created under Clause 1(i) is an integrated circuit which includes a CPU then either: (a) such CPU shall only be manufactured under licence from Arm; or (b) such CPU is neither substantially compliant with nor marketed as being compliant with the Arm instruction sets licensed by Arm from time to time;

(ii) the licences granted in Clause 1(iii) shall not extend to any portion or function of a product that is not itself compliant with part of the relevant AMBA Specification; and

(iii) no right is granted to LICENSEE to sublicense the rights granted to LICENSEE under this Agreement.

3. Except as specifically licensed in accordance with Clause 1, LICENSEE acquires no right, title or interest in any Arm technology or any intellectual property embodied therein. In no event shall the licences granted in accordance with Clause 1 be construed as granting LICENSEE, expressly or by implication, estoppel or otherwise, a licence to use any Arm technology except the relevant AMBA Specification.

4. THE RELEVANT AMBA SPECIFICATION IS PROVIDED "AS IS" WITH NO REPRESENTATION OR WARRANTIES EXPRESS, IMPLIED OR STATUTORY, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF SATISFACTORY QUALITY, MERCHANTABILITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE, OR THAT ANY USE OR IMPLEMENTATION OF SUCH ARM TECHNOLOGY WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADE SECRETS OR OTHER INTELLECTUAL PROPERTY RIGHTS.

5. NOTWITHSTANDING ANYTHING TO THE CONTRARY CONTAINED IN THIS AGREEMENT, TO THE FULLEST EXTENT PERMITTED BY LAW, THE MAXIMUM LIABILITY OF ARM IN AGGREGATE FOR ALL CLAIMS MADE AGAINST ARM, IN CONTRACT, TORT OR OTHERWISE, IN CONNECTION WITH THE SUBJECT MATTER OF THIS AGREEMENT (INCLUDING WITHOUT LIMITATION (I) LICENSEE'S USE OF THE ARM TECHNOLOGY; AND (II) THE IMPLEMENTATION OF THE ARM TECHNOLOGY IN ANY PRODUCT CREATED BY LICENSEE UNDER THIS AGREEMENT) SHALL NOT EXCEED THE FEES PAID (IF ANY) BY LICENSEE TO ARM UNDER THIS AGREEMENT. THE EXISTENCE OF MORE THAN ONE CLAIM OR SUIT WILL NOT ENLARGE OR EXTEND THE LIMIT. LICENSEE RELEASES ARM FROM ALL OBLIGATIONS, LIABILITY, CLAIMS OR DEMANDS IN EXCESS OF THIS LIMITATION.

6. No licence, express, implied or otherwise, is granted to LICENSEE, under the provisions of Clause 1, to use the Arm tradename, or AMBA trademark in connection with the relevant AMBA Specification or any products based thereon. Nothing in

Clause 1 shall be construed as authority for LICENSEE to make any representations on behalf of Arm in respect of the relevant AMBA Specification.

7. This Licence shall remain in force until terminated by you or by Arm. Without prejudice to any of its other rights if LICENSEE is in breach of any of the terms and conditions of this Licence then Arm may terminate this Licence immediately upon giving written notice to You. You may terminate this Licence at any time. Upon expiry or termination of this Licence by You or by Arm LICENSEE shall stop using the relevant AMBA Specification and destroy all copies of the relevant AMBA Specification in your possession together with all documentation and related materials. Upon expiry or termination of this Licence, the provisions of clauses 6 and 7 shall survive.

8. The validity, construction and performance of this Agreement shall be governed by English Law.

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by Arm and the party that Arm delivered this document to.

Product Status

The information in this document is final, that is for a developed product.

Web Address

<http://www.arm.com>

1 Introduction

This document lists errata on AMBA CHI Issue E.b.

Each errata description is organized as a brief reason for the change, along with the precise change.

1.1 Classification of the change

Each listed item has a classification ID, of the form XYYY, where:

X is the errata classification type as follows, C, R, E or D:

C	Clarification	Informative change only
R	Relaxation	Backward-compatible normative change, modifying existing functionality
E	Enhancement	Backward-compatible normative change, adding new functionality
D	Defect	Non-backward compatible normative change

YYY is an Arm internal tracking number.

2 New or Updated Errata

2.1 C597: Cancellation of CopyBack requests following overlapping snoop

Affects:

CHI-B, CHI-C, CHI-D, CHI-E.a, CHI-E.b

Description:

If a snoop occurs between the sending of a CopyBack request and sending of the associated CopyBackWrData, it is possible that a Dirty copy of that line may be passed. If the local line has been invalidated, then the Requester must send a CopyBackWrData_I response. If the line remains allocated in the Requester, it can either provide the data or cancel the transaction. CopyBackWrData_I indicates to the Home that the CopyBack is canceled.

CopyBackWrData provides an indication of the state of the cache at the RN-F when the data is sent, encoded on Resp[2:0]. When the Dirty responsibility has been taken away from a Requester by a snoop, it is possible that the line may still be allocated but a CopyBackWrData_I provided, which means that the cache state information in the response is imprecise and must be ignored.

The specification will be clarified in a number of places to make this clearer.

The precise change(s):

In section 4.11.1 (At the RN-F node) on page 4-242 of CHI-E.b, the following text:

- A Request Node is permitted to not send valid CopyBack Data, if the cache line state after the Snoop response is sent is I or SC. The cache state in the WriteData response, after CopyBack Data is taken away by the snoop, must be I and all byte enables must be deasserted and the corresponding data must be set to zero.

Will be updated to:

- If the cache line state after the Snoop response is sent is I, the cache state in the CopyBackWrData response must be I, all byte enables must be deasserted and the corresponding data must be set to zero.
- If the cache line state after the Snoop response is sent is UC or SC, a Request Node is permitted to not send valid CopyBack Data. If the RN decides not to send valid CopyBack data, the cache state in the CopyBackWrData response must be I, all byte enables must be deasserted and the corresponding data must be set to zero.

In Table 4-27 (Permitted WriteData responses and Opcode and Resp field encodings) on page 4-200 of CHI-E.b the following row:

Response	DAT Opcode	Resp[2:0]	Description
CopyBackWrData_I	0x2	0b000	Data corresponding to a CopyBack request. Cache line state when data was sent is I and the data in the response is not valid.

Will be updated to:

Response	DAT Opcode	Resp[2:0]	Description
CopyBackWrData_I	0x2	0b000	Indicates a CopyBack request has been canceled. The data in the response must be zero and all BE must be deasserted. The cache state in the response is imprecise and must be ignored.

In Table 4-39 (Requester cache state transitions for Write request transactions) on page 4-219 of CHI-E.b, the WriteCleanFull rows:

Request Type	Cache state at Requester			WriteData response	Comp response
	Initial	Before Write Data Response	Final		
WriteCleanFull	UD	UD	UC	CBWrData_UD_PD	CompDBIDResp
		UC	UC	CBWrData_UC	CompDBIDResp
		SD	SC	CBWrData_SD_PD	CompDBIDResp
	UD, SD	SC	SC	CBWrData_SC	CompDBIDResp
		I	I	CBWrData_I	CompDBIDResp

Will be updated to:

Request Type	Cache state at Requester			WriteData response	Comp response
	Initial	Before Write Data Response	Final		
WriteCleanFull	UD	UD	UC	CBWrData_UD_PD	CompDBIDResp
		UC	UC	CBWrData_UC	CompDBIDResp
		UC	UC	CBWrData_I	CompDBIDResp
	UD, SD	SD	SC	CBWrData_SD_PD	CompDBIDResp
		SC	SC	CBWrData_SC	CompDBIDResp
		SC	SC	CBWrData_I	CompDBIDResp
		I	I	CBWrData_I	CompDBIDResp

2.2 D638: ReadReceipt not optional in certain transaction flows

Affects:

CHI.E-b

Description:

For a Read transaction, Home is permitted to deallocate the request after receiving a ReadReceipt without waiting for a CompAck if the original request from the RN did not have an ordering requirement. The receiving of ReadReceipt by Home confirms that the Subordinate will not send a RetryAck response for that transaction.

In the transaction flow diagrams that were included as part of the CHI-E.b update, the allocating read diagram erroneously indicated that the ReadReceipt in the ReadNoSnpSep DMT flow was optional. It is in fact required as the sending of CompAck is only dependent on CompData or RespSepData, not DataSepResp. In a DMT flow, Home has no awareness of when DataSepResp is sent and so requires a ReadReceipt to confirm that a RetryAck response will not be seen.

In general, it can be stated that to use a DMT flow with a ReadNoSnpSep transaction from Home, a ReadReceipt must always be requested unless the following are all true for the original request from the RN:

- It was a non-allocating Read
- It had ExpCompAck asserted
- It had an ordering requirement

In this scenario, the RN must see RespSepData and at least one DataSepResp packet before it can send CompAck to Home. CompAck confirms to the Home that it will also not see a RetryAck for that transaction, meaning there is no need for a separate ReadReceipt response.

The precise change(s):

In Figure 2-1 "Allocating Read" on page 2-42, the ReadReceipt in flow 4 "Response from Home, Data from Subordinate", the optional box around the ReadReceipt message will be removed, making the sending of this message mandatory. To align with this, the following text in step 4 on the following page will be changed from:

Optionally, when the Home requests a ReadReceipt response, the Subordinate returns a read receipt, ReadReceipt, to the Home.

To:

The Subordinate returns a read receipt, ReadReceipt, to the Home.

The following text in step 4 "Response from Home, data from Subordinate" of the Non-Allocating Read flow on page 2-46 will also be changed from:

Optionally, when the Home requests a ReadReceipt response, the Subordinate returns a read receipt, ReadReceipt, to the Home. The Home must do this when a completion acknowledge is not required.

To:

Optionally, when the Home requests a ReadReceipt response, the Subordinate returns a read receipt, ReadReceipt, to the Home. The Home must request a ReadReceipt unless the original Request indicates a requirement for both ordering and a completion acknowledge.

Additionally Table 2-6 "Permitted DMT and DCT for ReadNoSnp and ReadOnce* from an RN" on page 2-48 will have the Notes column text for the row representing Order[1:0]=00, ExpCompAck=1, DMT=Y, DCT=Y changed from:

For DMT, Home can ensure the request to SN is not given a RetryAck response by either obtaining the Request Accepted response from SN or waiting for CompAck response.

To:

For DMT, to ensure that the request from Home to the SN is not given a RetryAck response:

- When using ReadNoSnp to the SN, Home must either obtain a ReadReceipt from the SN or wait for the CompAck response from the Request Node.
- When using ReadNoSnpSep to the SN, Home must obtain a ReadReceipt from the SN

2.3 C624: DVM payload encoding for Instruction cache invalidations

Affects:

CHI-E.a, CHI-E.b

Description:

To help clarify the field positioning in the DVMOp and SnpDVMOp request payloads, CHI-E.b had a new implementation of Table 8-8 vs CHI-E.a. The table wrongfully omitted the field locations of VI Valid and Virtual Index. The additional tables that detailed the various DVM operations also had their bit position values removed, as these differ between the DVMOp and SnpDVMOp request payloads due to the width of the Req.Addr and Snp.Addr fields. This combination of changes could cause confusion when establishing the encoding for Physical Instruction Cache Invalidate operations.

The precise change(s):

Table 8-8 will be modified in the following ways:

- Adding of "VI Valid[0]" to Req.Addr[5]/Dat.Data[5] row in the Request column where "VMID Valid" currently exists
- Adding of "VI Valid[0]" to Snp.Addr[2] row in the Part 1 column where "VMID Valid currently exists"
- Adding of "VI Valid[1]" to Req.Addr[6]/Dat.Data[6] row in the Request column where "ASID Valid" currently exists
- Adding of "VI Valid[1]" to Snp.Addr[3] row in the Part 1 column where "ASID Valid" currently exists
- Adding of "VI[27:20]" to Req.Addr[21:14]/Dat.Data[21:14] row in the Request column where "VMID[7:0]" currently exists
- Adding of "VI[27:20]" to Snp.Addr[18:11] row in the Request column where "VMID[7:0]" currently exists
- Adding of "VI[19:12]^d" to Req.Addr[37:22]/Dat.Data[37:22] row in the Request column where "ASID[15:0]" currently exists
- Adding of "VI[19:12]^d" to Snp.Addr[34:19] row in the Request column where "ASID[15:0]" currently exists

Additionally, the following footnote will be added to Table 8-8:

d. When used as Virtual Index (VI), Req.Addr[37:30], Dat.Data[37:30] and Snp.Addr[34:27] can take any value

Table 8-26 will have the column header "Virtual Index" changed to "VI Valid".

3 Previous Errata

3.1 C605: DBID values in Comp and DBIDResp messages that originate from difference sources in DWT flow have no relationship

Affects:

CHI-E.a, CHI-E.b

Description:

When a write transaction takes advantage of the DWT flow, the source of the DBIDResp and Comp response messages into the original requester will be different. DBIDResp will originate from the Subordinate and Comp will originate from the Home.

There is no requirement or relationship between the Comp.DBID value returned by the Home and the DBIDResp.DBID value returned by the Subordinate during DWT.

The precise change(s):

On page 2-90, in section 2.5.9 "Data Buffer ID" the following text:

A Comp response message sent separate from a DBIDResp or DBIDRespOrd message for a Write transaction must include the same DBID field value in the Comp and DBIDResp or DBIDRespOrd message.

Will be updated to:

A Comp response message sent separate from a DBIDResp or DBIDRespOrd message for a Write transaction must include the same DBID field value in the Comp and DBIDResp or DBIDRespOrd message when the two messages originate from the same source.

In addition, at the bottom page 2-106 in the "WriteNoSnp transaction" identifier field usage description, the following Note:

There is no ordering requirement between the separate DBIDResp and Comp responses. It is required that the values used are identical.

Will be updated to:

There is no ordering requirement between the separate DBIDResp and Comp responses. It is required that the values used are identical when the two messages originate from the same source.

3.2 C606: Data_Check and Check_Type property descriptions

Affects:

CHI-D, CHI-E.a, CHI-E.b

Description:

The Data_Check and Check_Type properties are used to declare the parity capabilities supported on an interface. There is an overlap between these two properties and additional clarity will be added to the specification to cover the case where they could potentially be viewed as being in contention. The parity support on an interface will be one of the following:

- No checking
- DataCheck field in the DAT packet
- Complete interface parity checking

The precise change(s):

On page 9-348 in section 9.6 "Data Check" the following line will change from:

The Data_Check property is used to indicate if Data Check is supported.

To:

The Data_Check and Check_Type properties are used to indicate if Data Check is supported in the DAT packet.

On pages 16-470 and 16-471 the Data_Check and Check_Type property descriptions in 16.1 "Interface properties and parameters" will be replaced with:

Data_Check

The Data Check property is used to indicate if the DataCheck field is present in the DAT packet.

- When not specified, or set to *False*, the DataCheck field is not present in the DAT packet unless specified by the Check_Type property.
- When set to *Odd_Parity*, Data Check is supported and the DataCheck field is present in the DAT packet. If Check_Type is defined and set to *Odd_Parity_Byte_All*, no DataCheck field is present in the DAT packet.

Check_Type

The Check Type property is used to indicate the protection scheme employed on an interface:

When not specified or set to *False*, no checking signals are present on the interface, unless specified by the Data_Check property.

- When set to *Odd_Parity_Byte_Data*, the DataCheck field is present in the DAT packet.
- When set to *Odd_Parity_Byte_All*, parity check signals are added to every channel. The signals added are detailed in 9.7.3 Interface parity check signals on page 9-351.

3.3 D607: Typographical error in Allocating Read figure

Affects:

CHI-E.b

Description:

The figure detailing the possible transaction flows for an Allocating Read Transaction is updated to correct a typographical error.

The precise change(s):

On page 2-42, Figure 2-1 "Allocating Read" will have SnpRespFwdwd replaced with SnpRespFwded

3.4 D609: Typographical error in RN to HN Write request attribute values table

Affects:

CHI-E.b

Description:

Table 4-13 lists the values permitted for key attributes in Write requests from Request Nodes to Home Nodes. It incorrectly lists WriteBackPtl twice, the second of which is in error and should be WriteCleanFull.

The precise change(s):

On page 4-178 in Table 4-13 "RN to HN Write request attribute values" replace:

WriteBackFull
WriteBackPtl

with:

WriteBackFull
WriteCleanFull

3.5 C612: Requirements for completer read response when MTE is not supported

Affects:

CHI-E.a, CHI-E.b

Description:

The Memory Tagging Extension (MTE) is a mechanism that is used to check the correct usage of data held in memory. In a given system, not all devices/addresses may support MTE. When a Read request is issued that also asks for MTE Tags (TagOp value of *Transfer* or *Fetch*), if the end device does not support MTE, the response must use a TagOp value of *Invalid*.

Section 12.11.3 "MTE not supported" can be interpreted as relaxing this constraint with its use of the term permitted, suggesting that there might be an alternative response that could also be given when this is not true.

The precise change(s):

On page 12-387 in section 12.11.3 "MTE not supported" the following text will change from:

When a Completer does not support MTE for the address in the request, then the Completer is permitted to send a TagOp value of *Invalid* in response to a Read transaction with a TagOp value of *Transfer* or *Fetch*.

To:

When a Completer does not support MTE for the address in the request, then the Completer must send a TagOp value of *Invalid* in response to a Read transaction with a TagOp value of *Transfer* or *Fetch*.

3.6 C621: Typographical error in example WriteUniqueStash with Data Pull transaction flow

Affects:

CHI-B, CHI-C, CHI-D, CHI-E.a, CHI-E.b

Description:

Figure 5-23 shows an example WriteUniqueStash with Data Pull transaction flow. The SnpResp provided by RN-F1 includes a Data Pull request, asking for the associated line to be returned to it. The HN-F returns CompData in a UniqueDirty state with PassDirty indicated. This means that RN-F1 must allocate the line in that state, but a label indicating this on the figure is currently missing.

The precise change(s):

At the arrow head signaling the return of CompData_UD_PD to RN-F1, a label of "I->UD" will be added to Figure 5-23.

3.7 C628: Requirements for ReadOnceMakeInvalid when invalidating a Dirty copy

Affects:

CHI-B, CHI-C, CHI-D, CHI-E.a, CHI-E.b

Description:

The completion of a ReadOnceMakeInvalid may result in a Dirty copy in the system being invalidated. This will be dependent on whether or not the Home accepts the invalidation hint included as part of the request. In order to ensure that the system remains coherent, if a Dirty copy is invalidated and not written back to memory, all other cached copies must be invalidated.

The precise change(s):

In section 4.2.1 on page 164 of CHI-E.b, the following paragraph will be updated from:

Read request to a Snoopable address region to obtain a snapshot of the coherent data. It is recommended, but not required, that all snoopable cached copies are invalidated. If a Dirty copy is invalidated, it does not need to be written back to memory.

To:

Read request to a Snoopable address region to obtain a snapshot of the coherent data. If a Dirty copy is invalidated, it does not need to be written back to memory. All cached copies must be invalidated if the invalidation hint is accepted and a Dirty copy is not written back to memory.