



NVIDIA T241 FABRIC-4

Errata

Document Summary

DE-11318-01

Issue	Summary	Fix Plan
FABRIC-4	In system configurations with 3 or more sockets, Arm GIC-700 may receive interleaved transfers of multi-transfer packets originating from GIC instances on different sockets.	Will not fix in hardware

Known Issues

FABRIC	
	In system configurations with three or more sockets, Arm® GIC-700 may receive interleaved transfers of multi-transfer packets originating from GIC instances on different sockets.
	<p>In some cases, inter-socket AXI Stream packets with multiple transfers may be interleaved by the fabric when presented to Arm Generic Interrupt Controller. GIC expects all transfers of a packet to be delivered without any interleaving.</p> <p>The following GICv3 commands may result in multiple transfer packets over inter-socket AXI Stream interface:</p> <ul style="list-style-type: none"> • Register reads from GICD_I* and GICD_N* • Register writes to 64-bit GICD registers other than GICD_IROUTER_n* • ITS command MOVALL <p>Multiple commands in GICv4+ utilize multiple transfer packets, including VMOVP, VMOVI, VMAPP, and 64-bit register accesses.</p>
FABRIC-4	<p>This issue impacts system configurations with three or more sockets that require multi-transfer packets to be sent over inter-socket AXI Stream interface between GIC instances on different sockets. GICv4 cannot be supported. The GICv3 software model can only be supported with the workaround described below.</p> <p>Single and dual-socket configurations are not impacted by this issue, and fully support GICv3 and GICv4.</p>
	<p>System configurations with three or more sockets will be restricted to supporting GICv3, with a further restriction on ITS command MOVALL.</p> <p>For correct register reads in GICv3, software must update the address used to access GICD on each socket rather than using the global address page. MMIO accesses must target the GIC on the socket that holds the data, for example SPI ranges owned by the socket's GIC. This ensures that the GIC will not utilize the inter-socket AXI Stream interface for servicing these GIC MMIO accesses.</p>
	<p>Fix plan</p> <p>Will not fix in hardware.</p>

Notice

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation ("NVIDIA") does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this specification.

NVIDIA products are not designed, authorized or warranted to be suitable for use in medical, military, aircraft, space or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this specification. NVIDIA does not accept any liability related to any default, damage, costs or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this specification, or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this specification. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA. Reproduction of information in this specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

Arm

Arm, AMBA and Arm Powered are registered trademarks of Arm Limited. Cortex, MPCore and Mali are trademarks of Arm Limited. All other brands or product names are the property of their respective holders. "ARM" is used to represent Arm Holdings plc; its operating company Arm Limited; and the regional subsidiaries Arm Inc.; Arm KK; Arm Korea Limited.; Arm Taiwan Limited; Arm France SAS; Arm Consulting (Shanghai) Co. Ltd.; Arm Germany GmbH; Arm Embedded Technologies Pvt. Ltd.; Arm Norway, AS and Arm Sweden AB.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2023 NVIDIA CORPORATION & AFFILIATES. All rights reserved.