



UNBREAKABLE-LINK® Adapter and Switch Technology

Table of contents

LLR Mode

LLR Negotiation

LLR Status

UNBREAKABLE-LINK® Switch Commands

show interfaces ib llr

NVIDIA adapter, switch, and interconnect products support a unique UNBREAKABLE-LINK® technology that ensures the network links stay connected (operational) even in a bad connectivity environment.

MLNX-OS offers PHY profile configuration for InfiniBand interfaces. PHY profile includes Link Level Retransmission (LLR) configuration. A PHY profile is bound to any InfiniBand interface.

Link Level Retransmission (LLR) is used on signal integrity marginal systems to decrease and/or eliminate the impact of physical errors on the system's performance.

LLR transmitter breaks the transmitted Layer 2 data stream into Cells and adds a CRC checksum to each cell.

LLR receiver checks the Cell CRC, in case there is no CRC errors, it forwards the cell and acknowledges the peer.

If a cell is dropped by the receiver the transmitter retransmits the cell.

Note

LLR is a NVIDIA proprietary feature and will only work with NVIDIA-to-NVIDIA ports.

Note

LLR is not operational for cables longer than 30m.

LLR Mode

The following LLR modes are applicable per port per speed:

- disable—no LLR

- enable—the port becomes passive, only if it got a request to use LLR it activates, otherwise it remains disabled
- enable-request—the port becomes active, it keeps sending LLR requests to the peer

LLR Negotiation

Both ports on the link perform LLR discovery and negotiation. In order the LLR to be in active state on the link, the following should apply:

- One port must be configured with LLR “enable-request” on the specified speed.
- The other port (peer) may be configured with LLR “enable-request” or “enable” on the same specified speed

Note

If both the local port and remote port configured with LLR “enabled” the LLR negotiation will not be activated—the ports will remain in LLR in-active state.

LLR Status

LLR status is a port parameter that states the current state of the LLR.

- Active—LLR is operationally running
- In-Active—LLR is not running

UNBREAKABLE-LINK® Switch Commands

show interfaces ib llr

	show interfaces ib [<number>] llr Displays LLR status	
Syntax Description	number	The interface number

Default	N/A
Configuration Mode	Any command mode
History	3.2.0500
Example	<pre> switch (config) # show interfaces ib llr ----- Interface LLR status ----- IB1/1 Inactive IB1/2 Inactive IB1/3 Inactive IB1/4 Inactive IB1/5 Inactive IB1/6 Inactive IB1/7 Inactive IB1/8 Inactive IB1/9 Inactive IB1/10 Inactive IB1/11 Inactive IB1/12 Inactive IB1/13 Inactive ... </pre>
Related Commands	
Notes	

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. 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 document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information 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 (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

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 this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform 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 document. NVIDIA accepts no 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 document 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 document. 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 document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND 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 NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA 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 NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. 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 Terms of Sale for the product.

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 2024, NVIDIA. PDF Generated on 11/18/2024