



# The SHIELD: Self-Healing Interconnect

## White Paper

# Table of Contents

1	Background.....	3
2	InfiniBand – the Network of Choice for High Performance Computing.....	3
3	Conclusion .....	5

---

# 1 Background

High performance computing (HPC) has always been used to solve complex problems. However, with the current trend of exponential data growth, new magnitudes of cluster computing scale are now necessary to tackle the computational challenges of today and tomorrow. Scientific computing can analyze ever larger and higher resolution models, while machine learning and artificial intelligence are increasingly becoming pervasive techniques on their own. Hybrid software techniques, which combine these methodologies, are creating new challenges for HPC applications.

In HPC, clusters have depended on a high-speed and reliable interconnect. Efficient inter-process communications depend on an interconnect fabric that is capable of high bandwidth and low latency while supporting a massive number of endpoints (compute and storage servers). MPI, SHMEM/PGAS, and UPC, fast access to large scale, shared storage, machine learning frameworks, and even new heterogeneous computing architectures all share common characteristics and requirements for a robust and resilient network in order to achieve maximum scalable performance.

---

## 2 InfiniBand – the Network of Choice for High Performance Computing

For many years, InfiniBand has been the network of choice for cluster computing, delivering services at a very high degree of reliability and performance. As InfiniBand networks continue to expand to accommodate higher scales of computation and storage capacity, the increase in adapters and switches, and in particular, the cables accompanying that growth, will occasionally sustain physical or electrical damage.

Traditional software mechanisms addressing this issue include job checkpointing, which creates a point-in-time snapshot of the computation. If the computation fails at a later point, the job will resume from the last successful state and point in time. Of course, network protocols have data integrity checks and retransmission mechanisms, but these methods all have a negative impact on performance and are impractical at very large scales.

In the case of today's networks, a traditional subnet manager will recognize failed links and recalculate routes to avoid the problem, but this can take up to 5 seconds for 1,000 nodes and 30 seconds for clusters with 10,000 or more endpoints – certainly not fast enough to ensure the seamless integrity of a running computation. In fact, no software mechanism can be responsive enough at very large scales to detect and fix fabrics that suffer from a link failure.

To address this problem, NVIDIA designed a new and innovative solution called SHIELD™ (Self-Healing Interconnect Enhancement for Intelligent Datacenters), which takes advantage of the intelligence already built into the latest generation of InfiniBand switches. By making the fabric capable with self-healing autonomy, the speed with which communications can be corrected in the face of a link failure can be sped up by 5000x, fast enough to save

communications from expensive retransmissions or absolute failure. SHIELD was first introduced in NVIDIA® Mellanox® Switch-IB® 2 EDR 100Gb/s switch devices, and is also supported in the NVIDIA Mellanox Quantum™ HDR 200Gb/s switch devices.

SHIELD supports two mechanisms for communications recovery. The first and simplest case is one in which a switch has more than one forward route to the desired destination. In this case, the switch can make an independent decision to forward the packet out an alternate port that sets it on the new viable route.




Figure 1. Simple Case: Local Fix

However, there are cases in which no alternative routes to the desired endpoint are available from the switch that experiences a failed link, such as a downstream switch in a Fat-Tree network topology. In this case, the switch can pass in-band information to another switch in the network, a switch that can select the most efficient alternative route so it can then take responsibility for rerouting the traffic. The total time required to perform this action is on the order of 1 microsecond, quick enough to allow communications that depend on those connections to continue successfully.

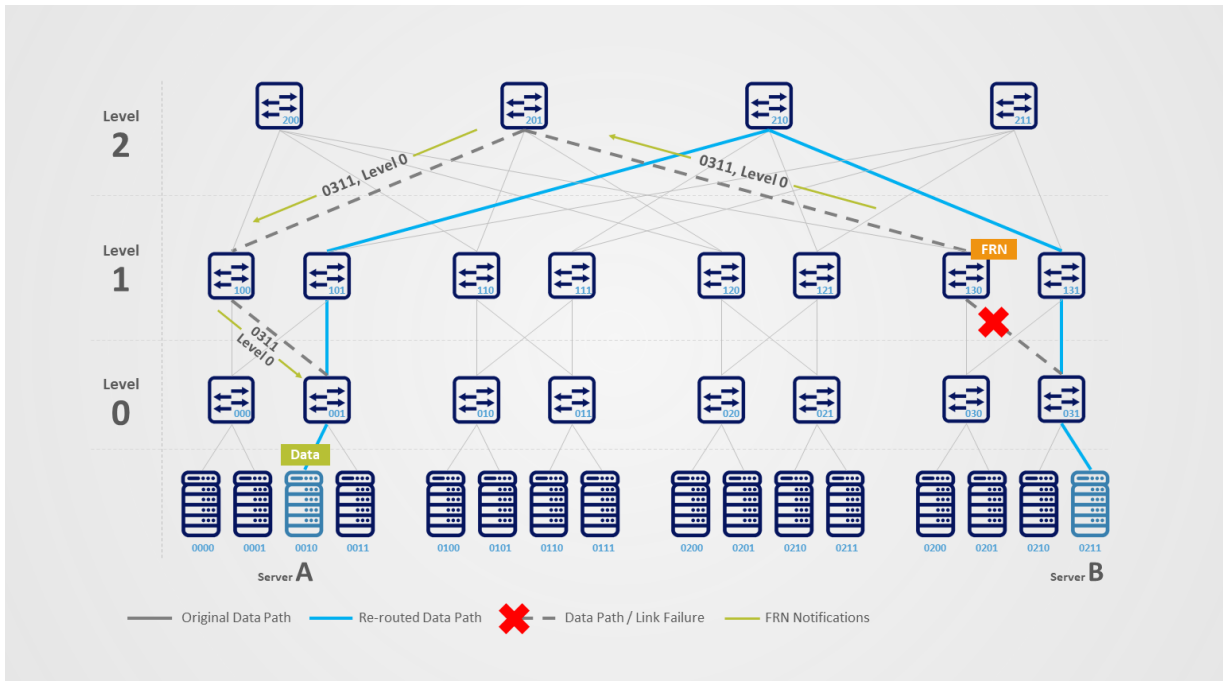


Figure 2. Remote Case: Using Fault Recovery Notifications

## 3 Conclusion

Networks have reached a greater degree of autonomy, detecting and correcting link failures and ensuring reliable, high performance data delivery, even as the industry approaches Exascale computing.

With SHIELD, NVIDIA Mellanox InfiniBand solutions are unleashing the possibilities of next generation high performance computing, today.

## 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. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries (collectively: "NVIDIA") make any 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, the NVIDIA logo, Mellanox, and SHIELD are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

## Copyright

© 2020 NVIDIA Corporation. All rights reserved.