



APPENDIX—Example of Migrating from FDR Gateways to NVIDIA Skyway

Table of contents

Overview

Procedure Description

Topology of the Migration Process

Traffic Flow

Migration Procedure

 Network Devices

Routes Directions

Overview

FDR SwitchX-2 based gateway (e.g. MSX6036G) is Proxy-ARP-based gateway. The Ethernet and InfiniBand networks are on the same IP subnet. NVIDIA Skyway (HDR/EDR gateway) is a Layer-3 (IP router) solution. The Ethernet and InfiniBand networks are not on the same IP subnet.

In order to migrate from MSX6036G-based gateway to NVIDIA Skyway, either the Ethernet or InfiniBand host needs to be moved to a different IP subnet. The following section describes the migration procedure, in case a customer wishes to handle this transition with minimum interruption to a working cluster.

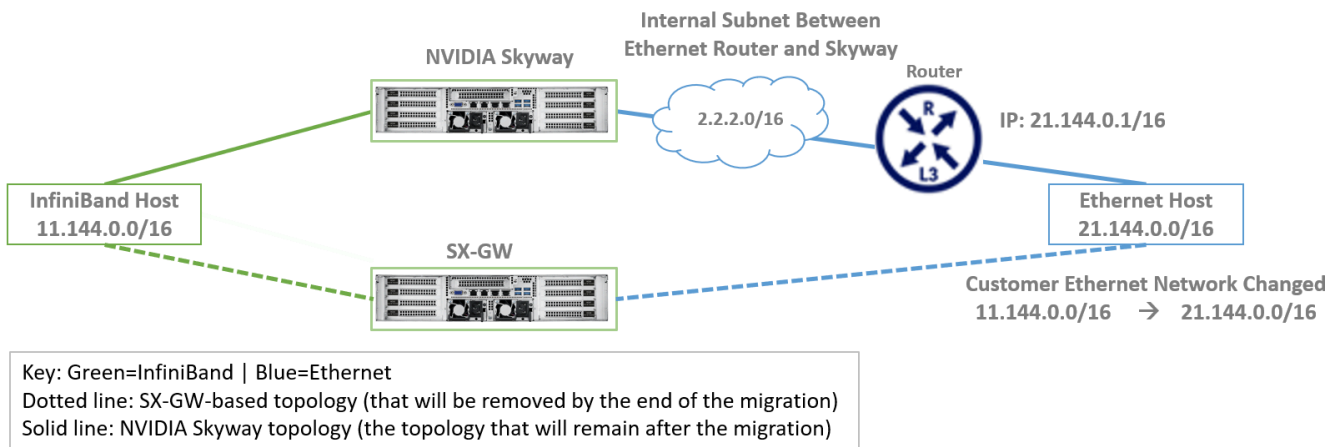
Note

Each operating system has its own configurations. In the examples below, Linux host Red Hat Enterprise Linux release 8.0 (Ootpa) and NVIDIA Onyx were used.

Procedure Description

In the following section, migration instructions are provided. The process assumes that the FDR network works well and the Skyway-based connectivity is gradually added. The diagram below depicts this process. The lower part of the diagram (with the dotted lines) shows the FDR-gateway connectivity. Gradually bring up the Skyway-based connectivity (the upper part of the diagram) by following the "Migration Procedure" below. At the end of the process, the traffic will flow through NVIDIA Skyway. During this process, short interruption in traffic flow will occur. After the completion of the migration, all traffic will go through the Skyway.

Topology of the Migration Process



Traffic Flow

- 1. Ethernet host to InfiniBand hosts with Iperf application.
- 2. Ethernet host to InfiniBand hosts with GPFS perf.

Migration Procedure

Network Devices

NVIDIA Skyway:

Step A: Connect NVIDIA Skyway to switches and configure IP addresses for NVIDIA Skyway InfiniBand and Ethernet interfaces (See "Starting the Command Line Interface (CLI)" in the [Getting Started](#) section)

Step B: Configure a static route on NVIDIA Skyway on the Ethernet network via the IP assigned on the Ethernet router's port (For more details, see "General Networkwide Deployment Guidelines" in the [Getting Started](#) section)

Ethernet Router:

Step C: Configure a router-port (via LAG or MLAG) on the Ethernet switch in the direction of NVIDIA Skyway.

Step D: Configure a route to the IPoIB network via the IP assigned to the gateway's Ethernet port-channel (for more details, see "General Networkwide Deployment Guidelines" in the [Getting Started](#) section)

Step E: Configure Interface VLAN to act as a "default gateway" (in the example, above, it is "21.144.0.1") for the Ethernet hosts (one example is that each port will be

connected to the server and access VLAN X).

****The traffic will be interrupted at this point!**

Ethernet Host:

Step F: Assign a new IP addresses and default gateway (in the example, above, the network was changed from "11.144.0.0/16" to "21.144.0.0/16") on the Ethernet hosts (can be done via DHCP or static).

****The traffic will resume after this point!**

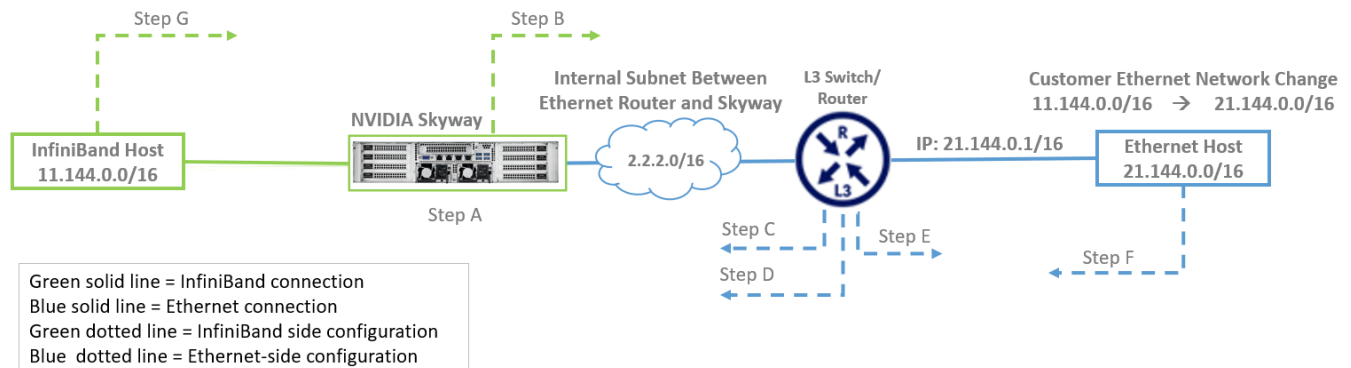
InfiniBand Host:

Step G: Assign default gateway on InfiniBand hosts (can be done via DHCP or static).

Note

This example shows modification of Ethernet subnet. It is equally optional to change the InfiniBand subnet instead.

Routes Directions



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 NONINFRINGEMENT, 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