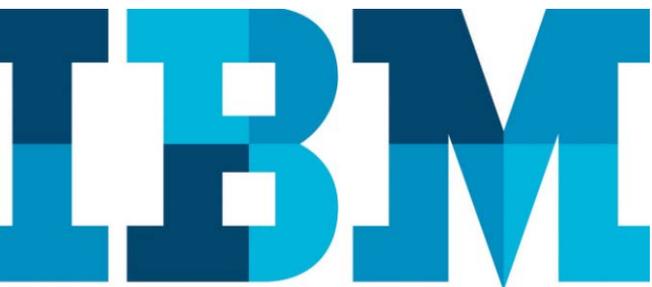


Configuring a Cavium adapter with Linux host attached configuration

Using system iSER target

Table of contents

<i>Getting started.....</i>	<i>2</i>
<i>Introduction</i>	<i>2</i>
<i>Advantages of Cavium iSER</i>	<i>3</i>
<i>Environment setup</i>	<i>3</i>
<i>Configuration on SVC</i>	<i>4</i>
<i>Cavium iSER initiator configuration on Linux hosts.....</i>	<i>4</i>
<i>Resources</i>	<i>10</i>
<i>About the authors.....</i>	<i>10</i>



Overview

Challenge

Setting up RoCE and iWARP profiles on Cavium adapters for iSER connectivity can be tricky.

Solution

Set up the RoCE and iWARP profiles on Cavium QL41262HLCU adapters and use a host with Cavium adapters for iSER host attach to IBM Spectrum Virtualize target.

The purpose of this paper is to provide detailed, step-by-step instructions on how to discover and establish connections from Cavium 25 Gb RDMA network interface card (RNIC). This paper demonstrates setting up both the Internet Wide Area RDMA Protocol (iWARP) and RDMA over Converged Ethernet (RoCE) profiles for Cavium adapters on Linux® hosts. The system configuration can be used for iSER host attach connections to IBM® System Storage® SAN Volume Controller (SVC) over high speed Ethernet links.

This paper is intended for people who are familiar with SVC, Internet Small Computer System Interface (iSCSI), and iSER but need a practical example to guide them through the steps of how to set the profile for CAVIUM adapter which supports both iWARP and RoCE protocols.

Getting started

IBM SVC and IBM Storwize® systems were being used as iSCSI targets till 8.1.0.0 release using 1 Gb and 10 Gb Ethernet ports. 8.1.1.0 release added support for 25 Gb RoCE and iWARP ports which could be used as iSCSI targets. This means that only iSCSI hosts could connect to IBM SVC or Storwize systems using the 25 Gb ports.

In the 8.2.0 release, IBM introduced support for iSER target on SVC and Storwize platforms using the 25 Gb RoCE and iWARP protocol. The RoCE adapter is **Mellanox ConnectX4-LX** supported on the host and the iWARP adapter is **Chelsio T6225-CR** supported on the host. The RoCE adapter can only establish iSER host-attached session to SVC target over the RoCE ports on the target. Similarly, the iWARP adapter can only establish iSER host attached session to the SVC/Storwize target over the iWARP ports on the target.

With the 8.3.0 release of SVC or Storwize, the QL41262HLCU Cavium adapter is also supported on the hosts. The ports on the Cavium adapters can be configured either as RoCE or iWARP ports.

Introduction

iSER is supported on Ethernet using two protocols – iWARP and RoCE. Most RNICs from different vendors support either the RoCE or iWARP protocols, not both. Users who are evaluating both these protocols either need to set up different servers, use multiple adapters, or physically replace the adapters when they need to switch from iWARP to RoCE connectivity or from RoCE to iWARP. Because this requires physical changes in the environment, it can be disruptive.

Advantages of Cavium iSER

iSER has the following advantages:

- The QL41262HLCU Cavium adapter is a second generation of standard compliant 25/50 Gb NIC that provides investment protection with support for concurrent RoCE, RoCEv2, and iWARP-universal RDMA.
- The adapter supports both iWARP and RoCE.
- The adapter allows the configuration to be switched by enabling the profile in the BIOS. Only change in BIOS settings and a reboot is required to enable the respective protocol.

Environment setup

This section describes the lab environment setup.

The setup includes two Linux® servers that has a 50 Gb logical unit number (LUN), allocated from an IBM Storwize storage controller (refer to Figure 1).

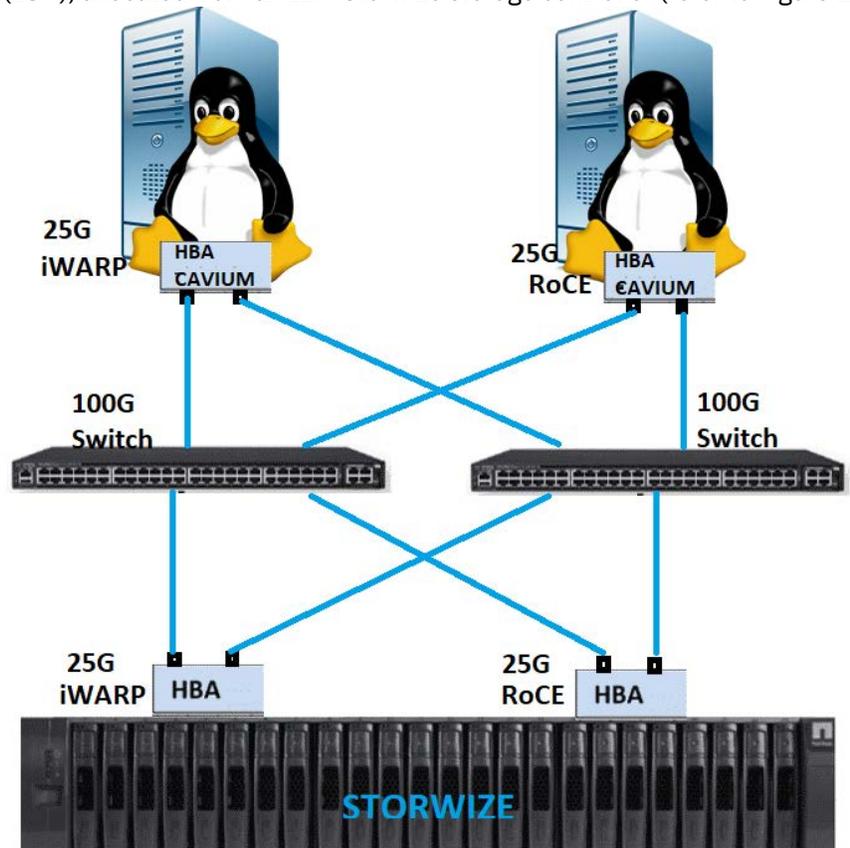


Figure 1. Lab setup with IBM Storwize as a storage controller

iSER initiator configuration is same on both SVC and Storwize systems. In the example used for this white paper, the Lab team used SVC as the iSER initiator. Same procedure is applicable for Storwize as well.

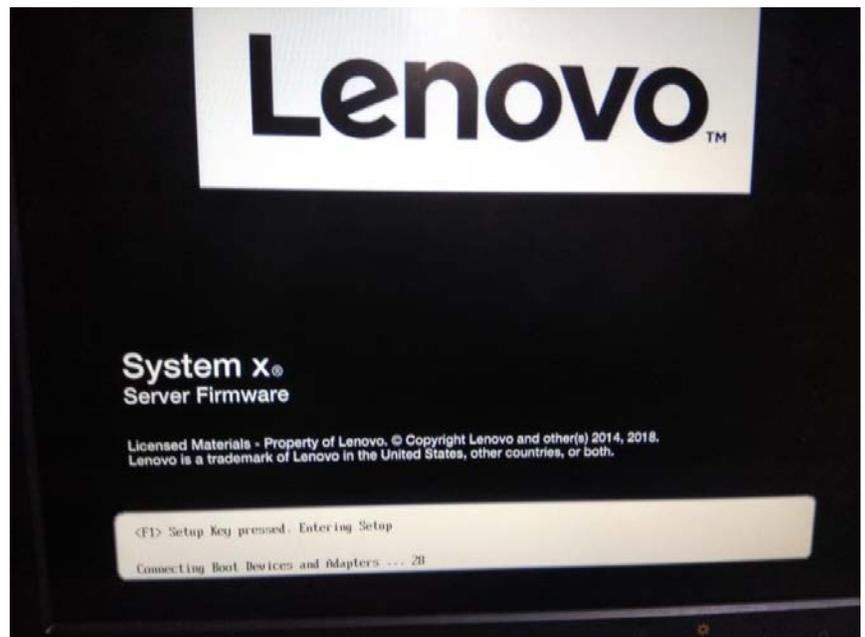
Configuration on SVC

Administrative model for configuring iSER is the same as iSCSI. iSER host-attached connectivity can be established in a few easy steps as explained at: <https://www.ibm.com/downloads/cas/NA9VAGXB>

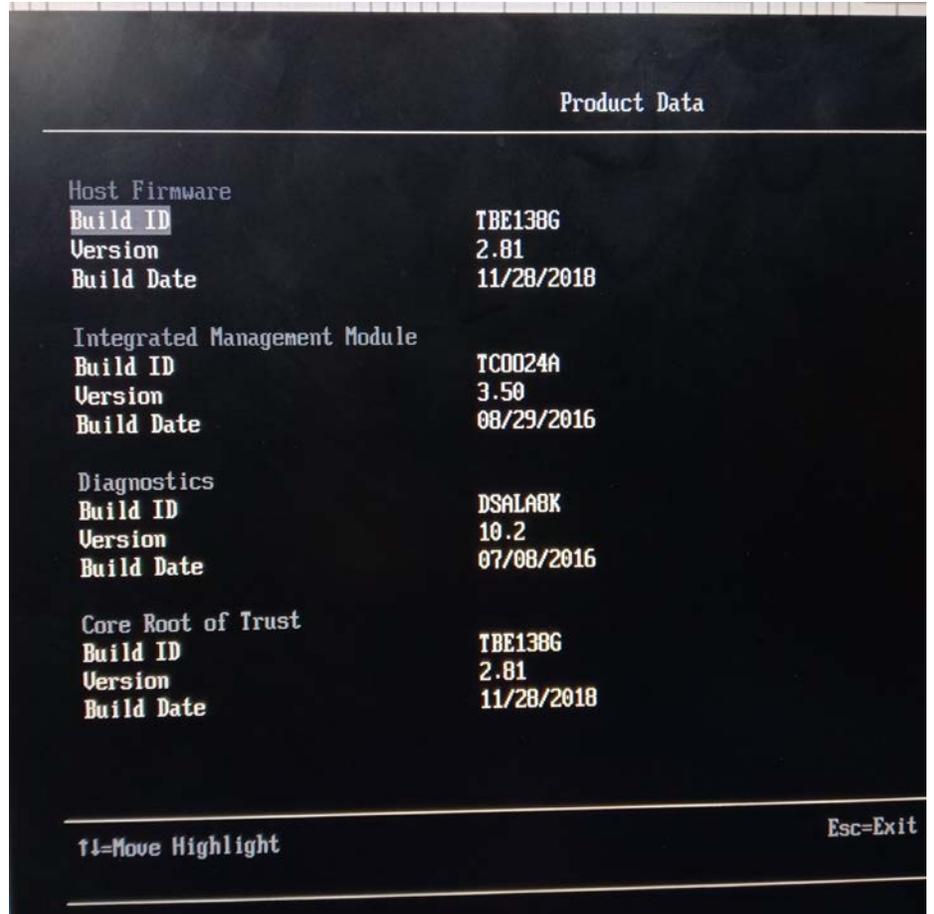
Cavium iSER initiator configuration on Linux hosts

To install the Linux software iSER initiator for iSER host attachment, set the iSCSI qualified name, target discovery, and authentication. Also, enable multipathing for the Linux hosts.

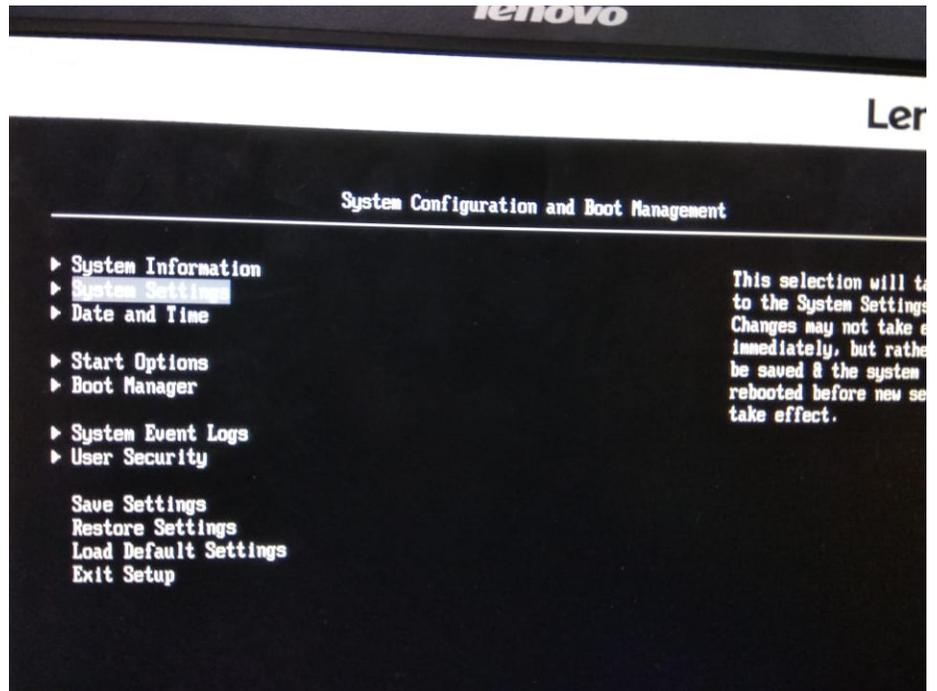
1. Ensure that the server to be configured has the latest BIOS version. You can follow any one of the two option: -
 - During bootup, verify the BIOS version
 - **Or**, enter BIOS while system bootup and press **F1**, go to setup, and follow the steps indicated on the screen.



2. Verify the version by checking the BIOS as shown in the following screen capture.



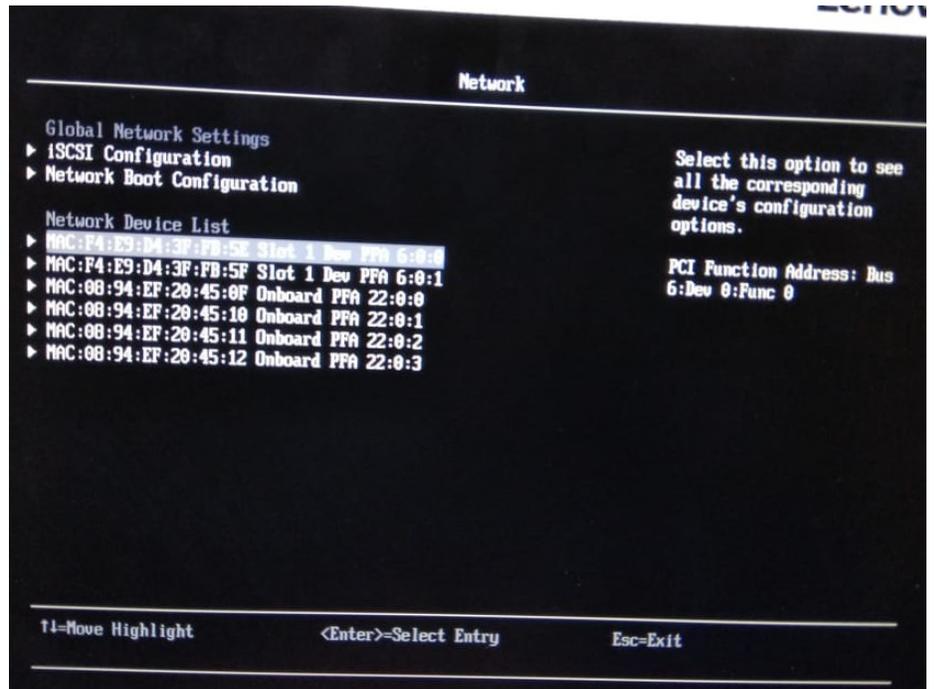
3. Go to System Configuration and Boot Management.



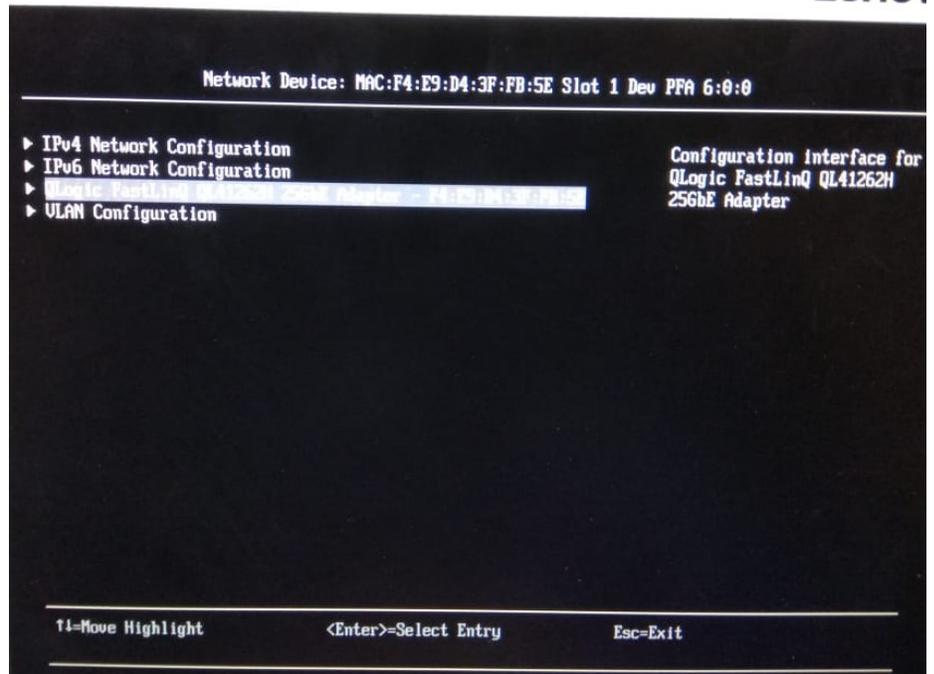
4. Select **System settings** and then select **Network**.



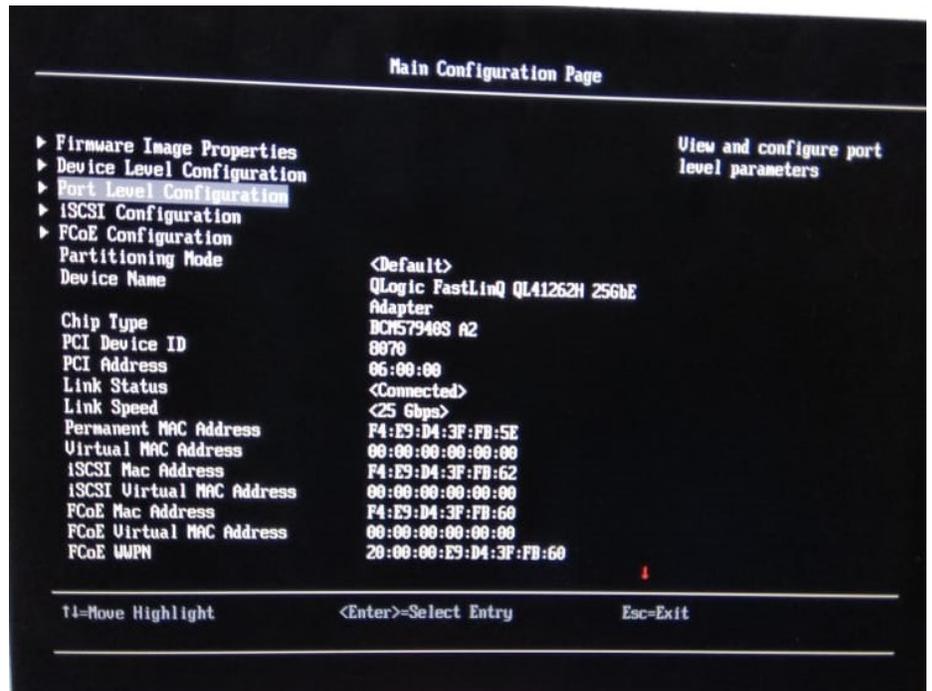
5. On the Network page, the network devices are displayed. In the list, verify the Cavium adapter ports. The ports will be listed if it is detected by BIOS. In this example, you can see that two ports are listed. Older BIOS would not list the Cavium adapter ports.



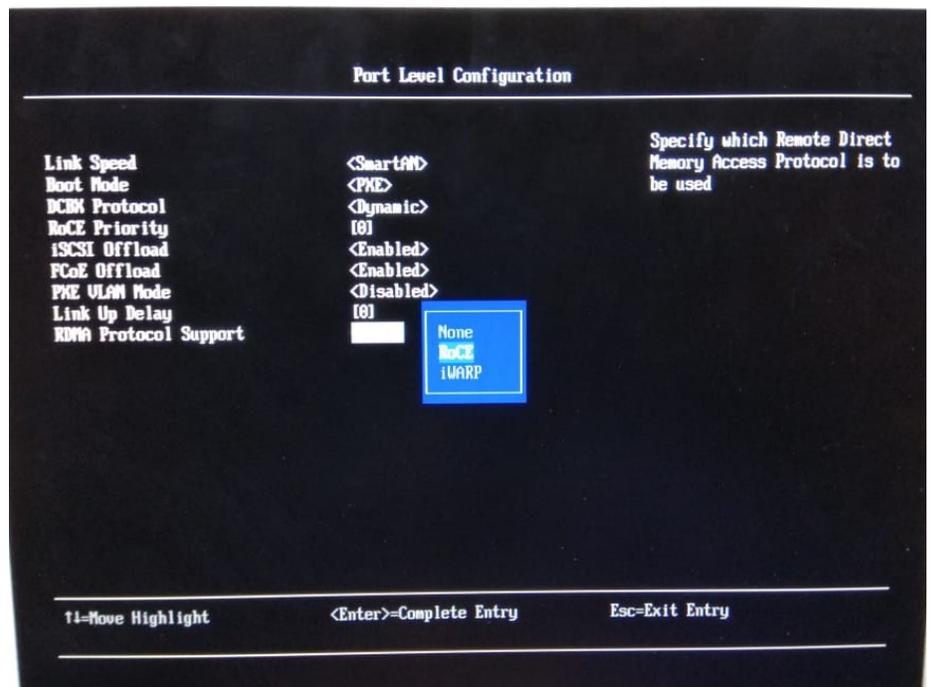
6. Enter the selected port and cross verify the Cavium model (as shown in the following screen capture).



7. Set the port-level profile configuration by selecting **QLogic FastLinQ QL41262HLCU 25GBE ADAPTER-<MAC>**.

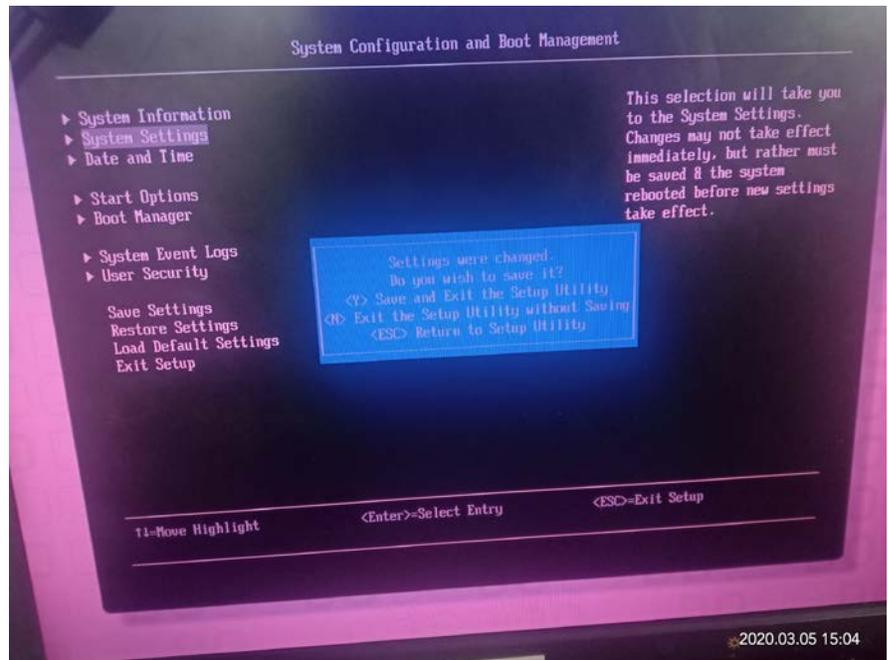


8. Select **Port Level Configuration**.



9. Select a protocol from the RDMA Protocol Support list. The options are RoCE, iWARP or NONE. (NONE stands for Native iSCSI support.)

10. Press the ESC key to exit the configuration screen. In the pop-up window that appears, type **Y** to save the configuration and exit BIOS.



11. After the system restarts, run the `lspci` command. In the output, notice that the adapter ports are displayed.

```
06:00.0 Ethernet controller: QLogic Corp. FastLinQ QL41000 Series 10/25/40/50GbE Controller (rev 02)
06:00.1 Ethernet controller: QLogic Corp. FastLinQ QL41000 Series 10/25/40/50GbE Controller (rev 02)
06:00.2 Ethernet controller: QLogic Corp. FastLinQ QL41000 Series 10/25/40/50GbE Controller (FCoE) (rev 02)
06:00.3 Ethernet controller: QLogic Corp. FastLinQ QL41000 Series 10/25/40/50GbE Controller (FCoE) (rev 02)
06:00.4 Ethernet controller: QLogic Corp. FastLinQ QL41000 Series 10/25/40/50GbE Controller (iSCSI) (rev 02)
06:00.5 Ethernet controller: QLogic Corp. FastLinQ QL41000 Series 10/25/40/50GbE Controller (iSCSI) (rev 02)
```

Resources

The following websites provide useful references to supplement the information contained in this paper:

- iSER host attachment from Linux hosts to IBM SVC and IBM Storwize family controllers
<https://www.ibm.com/downloads/cas/NA9VAGXB>
- iSER host attachment from Microsoft Windows servers to IBM SVC/ IBM Storwize family controllers
<https://www.ibm.com/downloads/cas/DP2BJ2ZJ>
- iSER host attachment from VMware hosts to IBM SVC and IBM Storwize/ FlashSystem controllers
<https://www.ibm.com/downloads/cas/32AVMKGQ>
- IBM Storwize V7000 Support page
[https://www.ibm.com/support/home/product/5402112/IBM_Storwize_V7000_\(2076\)](https://www.ibm.com/support/home/product/5402112/IBM_Storwize_V7000_(2076))
- IBM Storwize V7000 Supported hardware list
<https://www.ibm.com/support/pages/node/688673>
- IBM Storwize V7000 Knowledge Center
<https://www.ibm.com/support/knowledgecenter/ST3FR7>
- IBM SVC Knowledge Center
<https://www.ibm.com/support/knowledgecenter/STPVGU>
- IBM Redbooks®
<http://www.redbooks.ibm.com/redbooks.nsf/blueprints?OpenView>

About the authors

Rakesh Kumar Chinhara is a software engineer in the SVC team at IBM Systems Labs India with 10 years of IT experience. He holds a bachelor's degree in Electrical Engineering from the Biju Patnaik University of Technology . You can reach Rakesh at rchinhar@in.ibm.com.

Anuj Chandra is a software architect in the SVC team at IBM Systems Labs India with 20 years of IT experience. He holds a bachelor's degree in Computer Science from the University of Pune. You can reach Anuj at anujchan@in.ibm.com.



© Copyright IBM Corporation 2020
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of the Internal Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.



Please recycle