

# Intel<sup>®</sup> Ethernet FCoE

Quick Connect Guide (Linux\*)

---

Networking Division (ND)

*June 2014*

Revision 1.0



## LEGAL

---

By using this document, in addition to any agreements you have with Intel, you accept the terms set forth below.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>.

Intel and Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\* Other names and brands may be claimed as the property of others.

Copyright © 2014, Intel Corporation. All Rights Reserved.




## Revision History

---

Revision	Date	Comments
1.0	June 11, 2014	Initial public release



**NOTE:**      *This page intentionally left blank.*



## Contents

---

<b>1.0</b>	<b>Introduction</b> .....	1
<b>2.0</b>	<b>FCoE Basics</b> .....	1
<b>3.0</b>	<b>Administrative Ownership</b> .....	2
<b>4.0</b>	<b>Loading the FCoE Driver and Obtaining the Port (PWWN) and Node (NWWN) World Wide Names in Linux</b> .....	2
<b>5.0</b>	<b>Setting Up the Network</b> .....	6
<b>6.0</b>	<b>Setting Up the Storage</b> .....	8
<b>7.0</b>	<b>Setting the Remote Storage Disk in Linux</b> .....	10
<b>8.0</b>	<b>Summary</b> .....	11



**NOTE:**      *This page intentionally left blank.*



## 1.0 Introduction

This document is an introduction to FCoE storage connectivity from a Linux Server Administrator's perspective, and shows the basic connection from the Linux operating system to an FCoE storage target.

The goal of this document is to show the connection from a host perspective and the requirements provided by and to the Storage and Network Administrators. It is part of a series of FCoE Quick Connect guides for multiple operating systems.

To better understand the entire process, some limited basic port, zone, and storage grouping configurations are shown from Switch and Storage Administrator perspectives. The information in this document is for experienced System Administrators who are familiar with server, network, data center, and SAN storage concepts and technologies.

## 2.0 FCoE Basics

Fibre Channel over Ethernet (FCoE) is a computer technology that encapsulates Fibre Channel frames over an Ethernet network.

Fibre Channel is a set of standards that define reliable, high-speed transport for storage. Fibre Channel does not follow the OSI model, but rather utilizes its own five layer model: layer FC0-FC04. Interconnectivity is much like its OSI counterpart, composed of physical, Data Link, Network, Common Services, and Protocol Mapping layers. The FCoE protocol replaces the FC0 and FC1 layers of the Fibre Channel stack with the Physical and Data Link layers of Ethernet. Using Ethernet provides more flexibility and cost reductions to data center usages.

Some notable configuration parameters and conceptual terms for FCoE are: **Targets, Zones, and World-Wide Names.**

- Targets (or LUNs) are storage partitions assigned to a server from a remote storage array.
- Storage arrays use unique identifiers called World-Wide Names (WWNs) as an addressing schema to access these targets.
- Zoning or zones are the partitioning of a Fibre Channel or FCoE fabric into smaller subsets to restrict interference, add security, and to simplify management.

**Note:** Refer to [Table 1](#) to understand the Linux command line syntax that will be used for host configuration.

**Table 1. Command Reference**

Command	Description
chkconfig	Updates and queries run-level information for system services.
lldpad	Link Layer Discovery Protocol (LLDP) agent daemon.
dcbtool	Query and set the DCB settings of a DCB capable Ethernet interfaces.
fcoeadm	FCoE management tool for the Linux systems.



### 3.0 Administrative Ownership

Basic FCoE connectivity touches three technology disciplines: server, network, and storage.

The Server Administrator identifies and provides the WWN to both the Network and Storage Administrators. The Network Administrator ensures the network is setup end-to-end, and must provide configurations changes to the FCF-capable switch port and add the WWN to the FCoE fabric's zone. The Storage Administrator creates the LUN and host entity, then assigns each to a storage group to create the LUN masking.

Table 2 lists the responsibilities of the individual technology administrators.

**Table 2. Administrative Ownership**

Owner	Responsibilities
Server Administrator	<ol style="list-style-type: none"> <li>1. Turn on the FCoE service.</li> <li>2. Identify the host's Port/Node World-Wide Name.</li> <li>3. Provide the host's Port/Node World-Wide Name to the Network and Storage Administrators.</li> </ol>
Network Administrator	<ol style="list-style-type: none"> <li>1. Ensure end-to-end L2 connectivity.</li> <li>2. Configure Ethernet port with an FCoE-capable VLAN.</li> <li>3. Add the Port World Wide Name to the appropriate Zone.</li> </ol>
Storage Administrator	<ol style="list-style-type: none"> <li>1. Create server's host initiator.</li> <li>2. Create LUN Target.</li> <li>3. Create "Storage Group": <ol style="list-style-type: none"> <li>a. Assign host entity to the Storage Group.</li> <li>b. Assign target LUN to the Storage Group.</li> </ol> </li> </ol>

### 4.0 Loading the FCoE Driver and Obtaining the Port (PWWN) and Node (NWWN) World Wide Names in Linux

This section shows how to load the Intel® Ethernet DCB application tools for FCoE, and obtain the WWN in the Linux operating system.

1. There are several ways to load the FCoE tools onto Linux. If the tools were not loaded with the Operating System, they can be added post-installation. With the OS installation disk installed in the host and from a terminal window, run the following commands to install the tools:
  - a. **Red Hat Enterprise Linux:** Install the FCoE package with `yum -y install fcoe`, as shown in Figure 1.

**Note:** The package name is *fcoe* in RHEL 6.4 and below. The package name in RHEL 6.5 going forward is *fcoe-util*.

```
[root@rhel62-quad-b ~]# yum -y install fcoe*
Loaded plugins: product-id, refresh-packagekit, security, subscription-manager
Updating certificate-based repositories.
Setting up Install Process
Nothing to do
[root@rhel62-quad-b ~]#
```

**Figure 1. Installing on RHEL**





- b. **SUSE Linux Enterprise Server:** Install the FCoE package with `zypper install open-fcoe`, as shown in Figure 2.

**Note:** The package name is `fcoe` in SLES 11 and below. The package name in SLES 12 going forward is `fcoe-util`.

```
net08:~ # zypper install open-fcoe
Loading repository data...
Reading installed packages...
'open-fcoe' is already installed.
No update candidate for 'open-fcoe-1.0.26-0.17.4.x86_64'. The highest available
version is already installed.
Resolving package dependencies...

Nothing to do.
net08:~ #
```

Figure 2. Installing on SLES

2. From a terminal console, use the `ethtool` to identify the adapter to be used for FCoE, as shown in Figure 3.

```
[root@rhel62-quad-b ~]# ethtool eth2
Settings for eth2:
    Supported ports: [ FIBRE ]
    Supported link modes:   10000baseT/Full
    Supports auto-negotiation: No
    Advertised link modes:  10000baseT/Full
    Advertised pause frame use: No
    Advertised auto-negotiation: No
    Speed: 10000Mb/s
    Duplex: Full
    Port: Direct Attach Copper
    PHYAD: 0
    Transceiver: external
    Auto-negotiation: off
    Supports Wake-on: d
    Wake-on: d
    Current message level: 0x00000007 (7)
    Link detected: yes
[root@rhel62-quad-b ~]#
```

Figure 3. Loading Driver/Obtaining Port and Node WWN - Step 2

3. Use `chkconfig` to verify whether the `lldpad` (tools) and FCoE services are running, as shown in Figure 4.

```
[root@rhel62-quad-b ~]# chkconfig lldpad on
[root@rhel62-quad-b ~]# chkconfig fcoe on
[root@rhel62-quad-b ~]#
```

Figure 4. Loading Driver/Obtaining Port and Node WWN - Step 3



4. Start the FCoE service as shown in [Figure 5](#).

```
[root@rhel62-quad-b ~]# service fcoe start
Starting FCoE initiator service:
[root@rhel62-quad-b ~]#
```

**Figure 5. Loading Driver/Obtaining Port and Node WWN - Step 4**

5. Use `dcbtool sc` (set config) to turn Data Center Bridging onto the adapter being used for FCoE, as shown in [Figure 6](#).

```
[root@rhel62-quad-b ~]# dcbtool sc eth2 dcb on
Command:      Set Config
Feature:      DCB State
Port:         eth2
Status:       Successful
[root@rhel62-quad-b ~]#
```

**Figure 6. Loading Driver/Obtaining Port and Node WWN - Step 5**

6. Use `dcbtool sc` (set config) to turn `app:0` (FCoE) onto the desired FCoE adapter, as shown in [Figure 7](#). The subtype can be 0 or `fcoe`.

Subtypes:

- 0: `fcoe` - Fiber Channel over Ethernet (FCoE)
- 1: `iscsi` - Internet Small Computer System Interface (iSCSI)
- 2: `fip` - FCoE Initialization Protocol (FIP)

```
[root@rhel62-quad-b ~]# dcbtool sc eth2 app:0 e:1
Command:      Set Config
Feature:      Application FCoE
Port:         eth2
Status:       Successful
[root@rhel62-quad-b ~]#
```

**Figure 7. Loading Driver/Obtaining Port and Node WWN - Step 6**

7. Optionally, PFC (Priority Flow Control) can be set with all default arguments (`e:1 / a:1 / w:1`), as shown in [Figure 8](#).

args can include:

- `e: <0|1>` - controls feature enable
- `a: <0|1>` - controls whether the feature is advertised via DCBX to the peer
- `w: <0|1>` - controls whether the feature is willing to change its operational configuration based on what is received from the peer



```
[root@rhel62-quad-b ~]# dcbtool sc eth2 pfc e:1 a:1 w:1
Command:      Set Config
Feature:      Priority Flow Control
Port:         eth2
Status:       Successful
[root@rhel62-quad-b ~]#
```

Figure 8. Loading Driver/Obtaining Port and Node WWN - Step 7

8. The last step to enable FCoE is to add the configuration file. A configuration script (`cfg-ethx`) was created in the directory `/etc/fcoe`. Copy the script to a new file named after the Ethernet adapter to be used for FCoE, as shown in Figure 9.

There are three configurable settings within the script:

- `fcoe_enable=yes`
- `dcb_required=yes`
- `auto_vlan=yes`

```
[root@rhel62-quad-b fcoe]# cd /etc/fcoe/
[root@rhel62-quad-b fcoe]# cp cfg-ethx cfg-eth2
[root@rhel62-quad-b fcoe]# ls
cfg-eth2  cfg-ethx  config
[root@rhel62-quad-b fcoe]#
```

Figure 9. Loading Driver/Obtaining Port and Node WWN - Step 8

9. Before verifying the switch and storage setting, make one last check to ensure the FCoE service is running, as shown in Figure 10.

**Note:** When the Intel adapter created the VLAN tag, it read and used the switch's FCoE VLAN (110).

**Note:** For SLES 11 SP3 the syntax is `service boot_fcoe status`.

```
[root@rhel62-quad-b ~]# service fcoe status
11786
/usr/sbin/fcoemon -- RUNNING, pid=11786
Created interfaces: eth2.110-fcoe
[root@rhel62-quad-b ~]#
```

Figure 10. Loading Driver/Obtaining Port and Node WWN - Step 9

10. Two notable commands prior to moving forward are a run of the command `fcoeadm -i` (interface) to obtain WWN and interface state information as shown in Figure 11, and `lldptool -tni eth2` (get neighbor interface) to query received neighbor information as shown in Figure 12. Here we can see switch name and the Ethernet port being used.

```
[root@rhel62-quad-b ~]# fcoeadm -i eth2.110-fcoe
Description:      82599EB 10-Gigabit SFI/SFP+ Network Connection
Revision:        01
Manufacturer:    Intel Corporation
Serial Number:   90E2BA206B80
Driver:          ixgbe 3.4.8-k
Number of Ports: 1

Symbolic Name:   fcoe v0.1 over eth2.110-fcoe
OS Device Name:  host6
Node Name:       0x100090E2BA206B82
Port Name:       0x200090E2BA206B82
FabricName:      0x200B547FEEC17401
Speed:           10 Gbit
Supported Speed: 10 Gbit
MaxFrameSize:   2112
FC-ID (Port ID): 0xCA0040
State:           Online

[root@rhel62-quad-b ~]#
```

Figure 11. Loading Driver/Obtaining Port and Node WWN - Step 10

```
[root@rhel62-quad-b ~]# lldptool -tni eth2
Chassis ID TLV
MAC: 54:7f:ee:c1:74:10
Port ID TLV
Local: Eth1/9
Time to Live TLV
120
Port Description TLV
Ethernet1/9
System Name TLV
Net5010
```

Figure 12. Loading Driver/Obtaining Port and Node WWN - Step 10

## 5.0 Setting Up the Network

The Network Administrator owns network switch port configuration and end-to-end connectivity between the storage array and the server. The overall configuration of a Fibre Channel Forwarder (FCF) switch is outside the parameters of this document, but a basic view of the key SAN administration configurations of the switch provides a valuable view of the FC/FCoE process.

**Note:** There are no configuration steps in this section, only `show` commands.

**Note:** For clarity and accuracy of the configuration, it is recommended to use the *Cisco\* Nexus\* 5000 Series NX-OS Software Configuration Guide*.



1. Use the global command `show flogi database` to verify FCoE VSAN details as shown in [Figure 13](#). The key information to note is the assigned VSAN number and PWWN for zone identification and configuration. This verifies the PWWN we saw in [Section 4.0](#), Step 10.

```
Net5010(config)#  
Net5010# show flogi database  
-----  
INTERFACE          VSAN    FCID          PORT NAME          NODE NAME  
-----  
fc2/1              11      0xca02ef      50:06:01:6c:3b:20:45:69 50:06:01:60:bb:20:45:69  
                    [EMC_960_FC]  
vfc5               12      0x8d0040      20:00:90:e2:ba:0a:23:da 10:00:90:e2:ba:0a:23:da  
vfc8               11      0xca0021      20:00:90:e2:ba:20:68:52 10:00:90:e2:ba:20:68:52  
                    [NET05_X520-1]  
vfc9               11      0xca0040      20:00:90:e2:ba:20:6b:82 10:00:90:e2:ba:20:6b:82  
                    [NET06_X520-1]  
vfc17              12      0x8d0000      50:06:01:6c:3e:a0:07:83 50:06:01:60:be:a0:07:83  
  
Total number of flogi = 5.  
Net5010#
```

**Figure 13. Setting Up the Network - Step 1**

2. Use the `show vlan fcoe` command to verify that Ethernet VLAN 110 and VSAN 11 are mapped together to encapsulate Fiber Channel packets onto an Ethernet, as shown in [Figure 14](#). VLAN 110 must be added to the Ethernet interface to participate in VSAN 11.

```
Net5010# show vlan fcoe  
  
Original VLAN ID      Translated VSAN ID      Association State  
-----  
110                   11                       Operational  
120                   12                       Operational  
130                   13                       Operational  
Net5010#
```

**Figure 14. Setting Up the Network - Step 2**

3. A view of the interface shows the port configured as an edge trunk as shown in [Figure 15](#). Note that in VLAN 110, the FCoE VLAN is allowed. An edge trunk or edge port is the Cisco Nexus version of Portfast. It allows a port to bypass many of the spanning tree states and go directly into a forwarding state. Additionally, VLAN 353 allows the port to be used as a traditional IP network which can be added to host through a VLAN tag.

```
Net5010# show run interface ethernet 1/9

!Command: show running-config interface Ethernet1/9
!Time: Fri Apr 18 12:10:48 2014

version 5.2(1)N1(4)

interface Ethernet1/9
 description V_110_Clarion_FCoE
 switchport mode trunk
 switchport trunk allowed vlan 110,353
 spanning-tree port type edge trunk

Net5010#
```

Figure 15. Setting Up the Network - Step 3

4. Use the `show zone vsan 11` command to verify that the zone associated with VSAN 11 has been populated with the PWWN of the host, as shown in Figure 16.

```
Net5010(config)# show zone vsan 11
zone name zone11 vsan 11
 pwwn 20:00:90:e2:ba:20:68:52 [NET05_X520-1]
 pwwn 50:06:01:6c:3b:20:45:69 [EMC_960_FC]
 pwwn 20:00:90:e2:ba:20:6b:82 [NET06_X520-1]
Net5010(config)#
```

Figure 16. Setting Up the Network - Step 4

## 6.0 Setting Up the Storage

The Storage Administrator will set up the server initiator record and the target. Full administration is beyond the scope of this document, but the storage group settings provide a helpful visual of target masking.

1. The Storage Administrator will require the Port/Node WWN to configure the target initiator, as shown in Figure 17. The WWN is used as a unique identifier to perform direct path assignment from the host initiator to the target initiator.

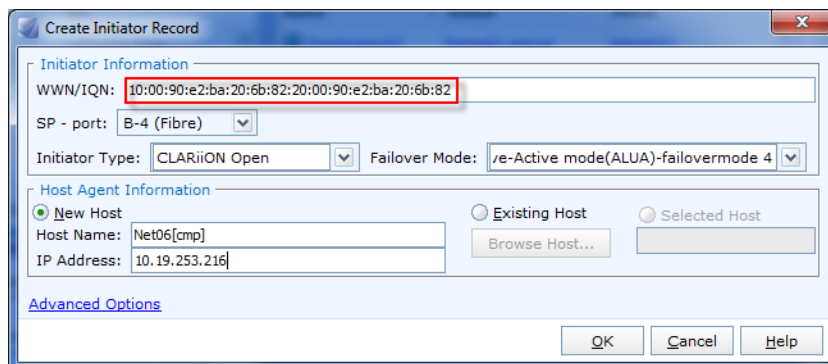


Figure 17. Setting Up the Storage - Step 1

2. Create a target LUN within a pool or volume within the storage array, as shown in Figure 18. In this case, a new 10 GB LUN is created in a RAID5 pool.




Figure 18. Setting Up the Storage - Step 2

3. The target LUN and the host initiator are mapped together in a storage group as shown in Figure 19.

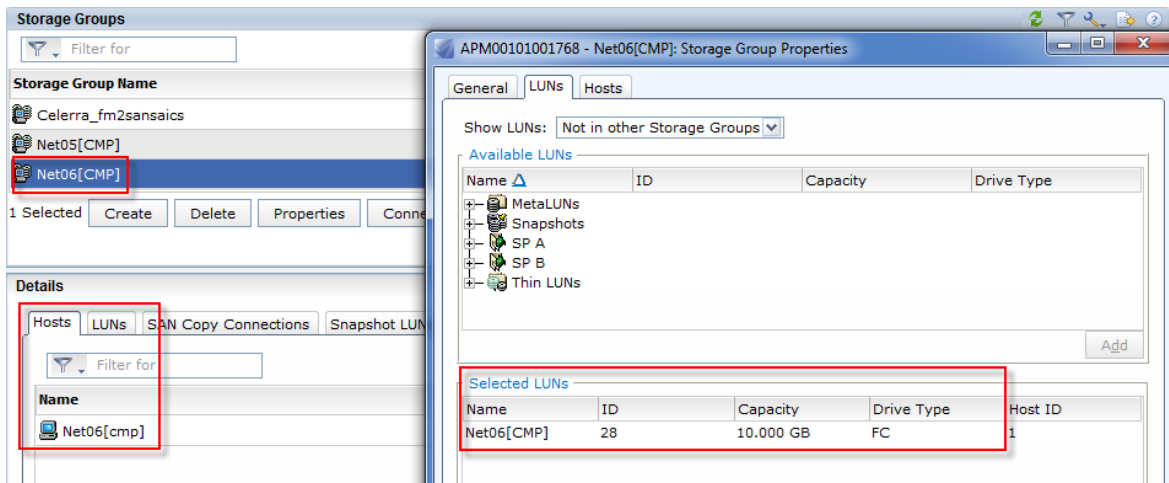


Figure 19. Setting Up the Storage - Step 3

## 7.0 Setting the Remote Storage Disk in Linux

This section shows the steps to view and manage the target LUN on the server once both the Storage and Network Administrators have created the storage target and the network has been fully configured.

1. With the switch and storage set up, verify that the host can see the target LUN. Run the command `fcoeadm -l` (list LUNs) to view interface, target and LUN information, as shown in [Figure 20](#).

```
[root@rhel62-quad-b ~]# fcoeadm -l
Interface:          eth2.110-fcoe
Roles:             FCP Target
Node Name:         0x50060160BB204569
Port Name:         0x5006016C3B204569
Target ID:         0
MaxFrameSize:     2048
OS Device Name:    rport-5:0-3
FC-ID (Port ID):  0xCA02EF
State:             Online

LUN #1 Information:
OS Device Name:    /dev/sdc
Description:       DGC VRAID (rev 0430)
Ethernet Port FCID: 0xCA0040
Target FCID:      0xCA02EF
Target ID:         0
LUN ID:           1
Capacity:          10.00 GB
Capacity in Blocks: 20971519
Block Size:        512 bytes
Status:            Attached

[root@rhel62-quad-b ~]#
```

**Figure 20. Setting the Remote Storage Disk in Linux - Step 1**

2. With the information found in Step 1, format the target LUN for use, as shown in [Figure 21](#). (`fdisk` is partition output only).

```
[root@rhel62-quad-b ~]# fdisk -b 4096 /dev/sdc
Disk /dev/sdc: 10.7 GB, 10737418240 bytes
64 heads, 32 sectors/track, 1280 cylinders, total 2621440 sectors
Command (m for help): u sector
Changing display/entry units to sectors

Command (m for help): p
Disk /dev/sdc: 10.7 GB, 10737418240 bytes
64 heads, 32 sectors/track, 1280 cylinders, total 2621440 sectors
Units = sectors of 1 * 4096 = 4096 bytes
Sector size (logical/physical): 4096 bytes / 4096 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x04a63d94
```

**Figure 21. Setting the Remote Storage Disk in Linux - Step 2**





## **8.0 Summary**

The Intel® Ethernet Converged Network Adapter Family offers some of the most flexible and scalable Ethernet adapters available for today's data centers. These adapters support Unified Networking and fully support a wide range of storage capabilities. Customers get the ease of Ethernet support along with the ability to converge a variety of applications in a single adapter.

For more configuration information on Intel® Server Adapters, go to:

<http://www.intel.com/support/network/sb/cs-009715.htm>

For more configuration information on Open-FCoE go to

<http://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/open-fibre-channel-over-ethernet-for-linux-based-ethernet-server-adapter-x520-brief.pdf>

<http://www.intel.com/content/dam/www/public/us/en/documents/guides/fcoe-user-guide.pdf>

[https://downloadcenter.intel.com/Detail\\_Desc.aspx?DwnldID=11848](https://downloadcenter.intel.com/Detail_Desc.aspx?DwnldID=11848)



**NOTE:**      *This page intentionally left blank.*