

Open FCoE for ESX* -based Intel® Ethernet Server X520 Family Adapters

Technical Brief v1.0

August 2011



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. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

IMPORTANT - PLEASE READ BEFORE INSTALLING OR USING INTEL® PRE-RELEASE PRODUCTS.

Please review the terms at http://www.intel.com/netcomms/prerelease_terms.htm carefully before using any Intel® pre-release product, including any evaluation, development or reference hardware and/or software product (collectively, "Pre-Release Product"). By using the Pre-Release Product, you indicate your acceptance of these terms, which constitute the agreement (the "Agreement") between you and Intel Corporation ("Intel"). In the event that you do not agree with any of these terms and conditions, do not use or install the Pre-Release Product and promptly return it unused to Intel.

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.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

Products 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.

Hyper-Threading Technology requires a computer system with an Intel® Pentium® 4 processor supporting HT Technology and a HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See http://www.intel.com/products/ht/Hyperthreading_more.htm for additional information.

Code names presented in this document are only for use by Intel to identify products, platforms, programs, services, etc. ("products") in development by Intel that have not been made commercially available to the public, i.e., announced, launched or shipped. They are never to be used as "commercial" names for products. Also, they are not intended to function as trademarks.


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 by visiting Intel's website at <http://www.intel.com>.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel Leap ahead., Intel Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside 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 © 2011, Intel Corporation. All Rights Reserved.



Contents

1.0	Introduction	5
2.0	Industry Support	5
2.1	Intel/EMC and Netapp Support Pages Available	5
3.0	Feature Introduction	5
4.0	Product Walk-Through	6



Revision History

Date	Revision	Description
August 2011	1.0	Initial release.



1.0 Introduction

In today's Data Center, virtualization has become the standard way of operating. Administrators are being constantly driven to lower operating costs while at the same time to provide better support for business-critical applications. Processor and platform advancements, together with changes in operating systems and applications has spurred data centers to quickly adopt 10 Gigabit Ethernet (10 GbE) as the de facto standard for interconnects. With the inherent I/O constraints associated with the Memory Controller Hub (MCH) removed in the latest generation chipsets (5500 series and later), applications that were not technically feasible a few short years ago are now becoming mainstream. Intel® Open FCoE for ESX*-based Intel X520 Adapters

One such application is Fiber Channel over Ethernet (FCoE). Fiber Channel over Ethernet is a standards-based solution that comprises standards from both FCIA (T11-FC-BB-5) as well IEEE* (Data Center Bridging). Historically, organizations have separated their LAN and SAN operations within a Data Center. With the advent of FCoE and DCB, The Information Trchnology (IT) group can now run both LAN, SAN and other traffic types over the same 10 GB infrastructure, resulting in lower costs: fewer cables/adapters and synergy among network and storage administration without a major upgrade to their existing SAN infrastructure.

2.0 Industry Support

Intel®, a long-time leader in the Data Center Ethernet market, has worked from the inception, with other industry leaders, to develop, test and ratify these standards. The Open FCoE solution, with Intel's X520 family of adapters listed below, has been certified with EMC* and Netapp*.

- Intel® Ethernet Server Adapter X520-DA2
- Intel® Ethernet Server Adapter X520-SR1
- Intel® Ethernet Server Adapter X520-SR2

2.1 Intel/EMC and Netapp Support Pages Available

<http://www.intel.com/support/network/adapter/pro100/sb/CS-032048.htm>


<http://www.intel.com/support/network/adapter/x520server/sb/CS-031565.htm>

3.0 Feature Introduction

With the release of VMWare's* ESX* 5.0, Intel's X520 family adapters are fully supported in the world's leading virtualization solutions. Intel worked closely with VMWare to integrate much of the Open FCoE stack--the solution developed by the community and embraced by Intel--into the ESX FCoE stack. The ESX hypervisor includes both FCoE and DCB natively and the Intel X520 family leverages this support.

Unlike other offloaded solutions that traditional HBA solutions have used, the software FCoE solution is fully integrated into the ESX hypervisor. This seamless integration provides assurances that other networking features within ESX, such as VLAN support, VMMotion, etc., and used in connection with FCoE, are architecturally robust with the native solution. It is worth noting that, although it is termed a software FCoE solution, that is a misnomer. There are indeed selective hardware FCoE offloads in the x520 including :

- Transmit and Receive (Tx - Rx) FC CRC Processing



- Receive Direct Data Placement (DDP)
- Receive Exchange-Id Packet Filtering
- Large Sequence Offload

These adapters offload the main FCoE data paths to improve throughput. Leveraging the full power of the current generation of Intel servers, together with these offload assists, performance with the Intel X520 adapters for real world workloads compares favorably with fully offloaded Fiber Channel host bus adapters within the ESX 5.0 environment.

4.0 Product Walk-Through

Included in the ESX 5.0 distribution are user level command line tools and the VSphere*/VCenter* graphical user interface (GUI) that make it easy for IT administrators to deploy FCoE. On a ESX host, users can discover, configure, enable, disable, or remove a given Intel X520 adapter as a FCoE adapter.

```
10.0.20.51 - PuTTY
~ # esxcfg-fcoe
No action provided
esxcfg-fcoe <action> [<options>]

Where <action> is one of:

-d|--discover=vmnicX [<options>]  Initiate FCoE adapter discovery on the given N
IC
-r|--remove-adapter=vmhbaXYZ      Destroy the specified FCoE adapter
-x|--deactivate-nic=vmnicW        Deactivate FCOE configuration for given NIC
-l|--list-vnports                  List discovered VNPorts associated with this h
ost
-N|--list-fcoe-nics                List FCoE-capable NICs with detailed informati
on
-n|--compact-list-fcoe-nics        List FCoE-capable NICs each on a single line,
with limited information
-e|--enable                        Enable an FCoE-capable NIC if it is disabled
-D|--disable                        Disable an FCoE-capable NIC if it is enabled (
requires
reboot to take effect)
-h|--help                          Show this message

And <options> are a set of:

-p|--priority={0-7}                Priority class to use for FCoE traffic
-v|--vlan=id                       VLAN ID to use for FCoE traffic
-a|--macaddress=xx:xx:xx:xx:xx:xx  MAC address to use for the underlying FCoE con
troller

Examples:

To discover FCoE adapters on a given NIC, using default settings
  esxcfg-fcoe -d vmnicX


To discover FCoE adapters on a given NIC, specifying only MAC address
  esxcfg-fcoe -d vmnicX -a MA

To discover FCoE adapters on a given NIC, specifying all settings
  esxcfg-fcoe -d vmnicX -p priority -v vlan -a MA

To remove an FCoE adapter
  esxcfg-fcoe -r vmhbaXYZ

To enable FCoE for a given NIC, specifying bandwidth and MAC address
  esxcfg-fcoe -e vmnicX -a MA

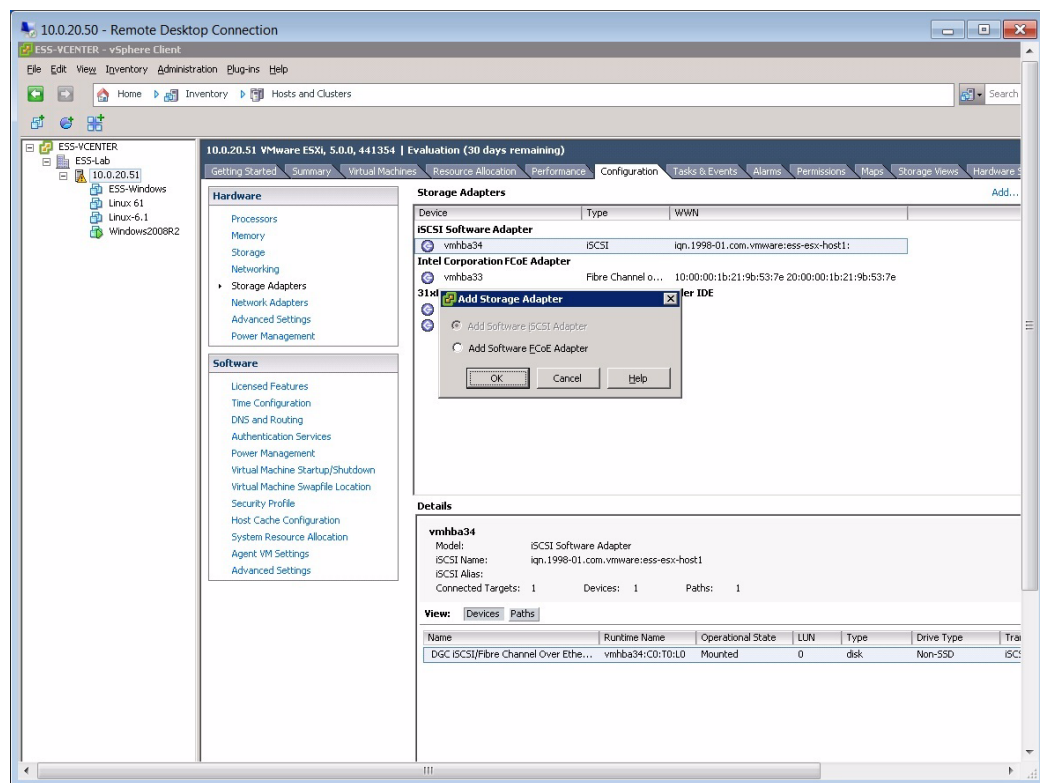
To disable FCoE for a given NIC
```



Users can configure the number of physical queues in the Intel X520 assigned to support FCoE in ESX 5.0 (the default is a single queue) as well as the queue depth on a lun via the host cli.

In ESX 5.0 VCenter, with a selected host highlighted and the Configuration tab selected, adding the Intel X520 adapter requires the following:

1. Under the Hardware menu, highlight Networking and use the Add Network Wizard to bind a VMKernel network connection to a vSphere vswitch that is connected to an Intel X520 Adapter. Please ensure that you also have a management connection to another vswitch in accordance with VMWare best practices.
2. Under the Hardware menu, highlight Storage adapter and add a Software FCoE adapter. This FCoE storage adapter instance is available for all guest OS'es.



For additional information, users are encouraged to consult the *ESX 5.0 Storage Administrators Guide* and the online help for configuration of FCoE adapters.

