Intel® Network Adapters User Guide

Restrictions and Disclaimers

Information for Intel® Boot Agent, Intel® Ethernet iSCSI Boot, Intel® FCoE/DCB can be found in the Re*mote Boot and Remote Storage for Intel® Ethernet Adapters and Devices* Guide.

Trademarks used in this text: *Dell* and the *DELL* logo are trademarks of Dell, Inc.; Intel is a trademark of Intel Corporation in the U.S. and other countries.

* Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Intel Corporation disclaims any proprietary interest in trademarks and trade names other than its own.

Restrictions and Disclaimers

The information contained in this document, including all instructions, cautions, and regulatory approvals and certifications, is provided by the supplier and has not been independently verified or tested by Dell. Dell cannot be responsible for damage caused as a result of either following or failing to follow these instructions.

All statements or claims regarding the properties, capabilities, speeds or qualifications of the part referenced in this document are made by the supplier and not by Dell. Dell specifically disclaims knowledge of the accuracy, completeness or substantiation for any such statements. All questions or comments relating to such statements or claims should be directed to the supplier.

Export Regulations

Customer acknowledges that these Products, which may include technology and software, are subject to the customs and export control laws and regulations of the United States (U.S.) and may also be subject to the customs and export laws and regulations of the country in which the Products are manufactured and/or received. Customer agrees to abide by those laws and regulations. Further, under U.S. law, the Products may not be sold, leased or otherwise transferred to restricted end users or to restricted countries. In addition, the Products may not be sold, leased or otherwise transferred to, or utilized by an end-user engaged in activities related to weapons of mass destruction, including without limitation, activities related to the design, development, production or use of nuclear weapons, materials, or facilities, missiles or the support of missile projects, and chemical or biological weapons.

14 March 2016

Overview

Welcome to the *User's Guide* for Intel® Ethernet Adapters and devices. This guide covers hardware and software installation, setup procedures, and troubleshooting tips for Intel network adapters, connections, and other devices.

Supported 40 Gigabit Network Adapters

- Intel® Ethernet 40G 2P XL710 QSFP+ rNDC
- Intel® Ethernet Converged Network Adapter XL710-Q2

NOTE: The total throughput supported by an Intel XL710-based adapter is 40 Gb/s, even when connected via two 40 Gb/s connections.

Supported 10 Gigabit Network Adapters

- Intel® Ethernet 10G 2P X520 Adapter
- Intel® Ethernet 10G X520 LOM
- Intel® Ethernet X520 10GbE Dual Port KX4-KR Mezz
- Intel® Ethernet 10G 2P X540-t Adapter
- Intel® Ethernet 10G 2P X550-t Adapter
- Intel® Ethernet 10G 4P X540/I350 rNDC
- Intel® Ethernet 10G 4P X520/I350 rNDC
- Intel® Ethernet 10G 2P X520-k bNDC
- Intel® Ethernet 10G 4P X710-k bNDC
- Intel® Ethernet 10G 2P X710-k bNDC
- Intel® Ethernet 10G X710-k bNDC
- Intel® Converged Network Adapter X710
- Intel® Ethernet 10G 4P X710/I350 rNDC
- Intel® Ethernet 10G 4P X710 SFP+ rNDC
- Intel® Ethernet 10G X710 rNDC

NOTE: The first port of an X710-based Adapter will display the correct branding string. All other ports on the same device will display a generic branding string.

Supported Gigabit Network Adapters and Devices

- Intel® Gigabit 2P I350-t Adapter
- Intel® Gigabit 4P I350-t Adapter
- Intel® Gigabit 4P I350 bNDC
- Intel® Gigabit 4P I350-t rNDC
- Intel® Gigabit 4P X540/I350 rNDC
- Intel® Gigabit 4P X520/I350 rNDC
- Intel® Gigabit 4P I350-t Mezz
- Intel® Gigabit 4P X710/I350 rNDC
- Intel® Gigabit 4P I350-t bNDC
- Intel® Ethernet Connection I354 1.0 GbE Backplane
- Intel® Gigabit 2P I350-t LOM
- Intel® Gigabit I350-t LOM
- Intel® Gigabit 2P I350 LOM

Installing the Network Adapter

If you are installing a network adapter, follow this procedure from step 1.
If you are upgrading the driver software, start with step 4.

1. Review system requirements.
2. Insert the PCI Express Adapter, Mezzanine Card, or Network Daughter Card into your server.
3. Carefully connect the network copper cable(s), fiber cable(s), or direct attach cables
4. Install the network drivers and other software
    - Windows Instructions
    - Linux Instructions

5. Testing the Adapter.

System Requirements

Hardware Compatibility

Before installing the adapter, check your system for the following minimum configuration requirements:

- IA-64-based (64-bit x86 compatible)
- One open PCI Express* slot (see the specifications of your card for slot compatibility)
- The latest BIOS for your system

Supported Operating Systems

Software and drivers are supported on the following 64-bit operating systems:

- Microsoft* Windows Server* 2012 R2
- Microsoft Windows Server 2012
- Microsoft Windows Server 2008 R2
- Microsoft Windows* 10
- Microsoft Windows 8.1
- Microsoft Windows 7
- VMWare* ESXi* 6.0 U2
- VMWare ESXi 5.5 U3
- Red Hat* Enterprise Linux* (RHEL) 7.2
- RHEL 7.1
- RHEL 6.7
- Novell* SUSE* Linux Enterprise Server (SLES) 12 SP1
- SLES 11 SP4

Cabling Requirements

Intel Gigabit Adapters

- 1000BASE-SX on 850 nanometer optical fiber:
  - Utilizing 50 micron multimode, length is 550 meters max.
  - Utilizing 62.5 micron multimode, length is 275 meters max.
- 1000BASE-T or 100BASE-TX on Category 5 or Category 5e wiring, twisted 4-pair copper:
  - Make sure you use Category 5 cabling that complies with the TIA-568 wiring specification. For more information on this specification, see the Telecommunications Industry Association's web site: www.-tiaonline.org.
  - Length is 100 meters max.
  - Category 3 wiring supports only 10 Mbps.

Intel 10 Gigabit Adapters

- 10GBASE-SR/LC on 850 nanometer optical fiber:
  - Utilizing 50 micron multimode, length is 300 meters max.
  - Utilizing 62.5 micron multimode, length is 33 meters max.
- 10GBASE-T on Category 6, Category 6a, or Category 7 wiring, twisted 4-pair copper:
  - Length is 55 meters max for Category 6.
  - Length is 100 meters max for Category 6a.
  - Length is 100 meters max for Category 7.
- 10 Gigabit Ethernet over SFP+ Direct Attached Cable (Twinaxial)
  - Length is 10 meters max.

Intel 40 Gigabit Adapters

- 40GBASE-SR/LC on 850 nanometer optical fiber:
  - Utilizing 50 micron multimode, length is 300 meters max.
  - Utilizing 62.5 micron multimode, length is 33 meters max.
- 40 Gigabit Ethernet over SFP+ Direct Attached Cable (Twinaxial)
  - Length is 7 meters max

OS Updates

Some features require specific versions of an operating system. You can find more information in the sections that describe those features. You can download the necessary software patches from support sites, as listed here:

- Microsoft Windows Server Service Packs: support.microsoft.com
- Red Hat Linux: www.redhat.com
- SUSE Linux: http://www.novell.com/linux/suse/
- ESX: http://www.vmware.com/

Ethernet MAC Addresses

Single-Port Adapters

The MAC address should be printed on a label on the card.

Multi-Port Adapters

Multiport adapters have more than one MAC address. The address for the first port (port A or 1) is printed on a label on the card.

Intel® Network Adapters Quick Installation Guide

Install the Intel PCI Express Adapter

1. Turn off the computer and unplug the power cord.
2. Remove the computer cover and the adapter slot cover from the slot that matches your adapter.
3. Insert the adapter edge connector into the PCI Express slot and secure the bracket to the chassis.
4. Replace the computer cover, then plug in the power cord.

NOTE: For information on identifying PCI Express slots that support your adapters, see your Dell system guide.

Attach the Network Cable

1. Attach the network connector.
2. Attach the other end of the cable to the compatible link partner.
3. Start your computer and follow the driver installation instructions for your operating system.

Install the Drivers

Windows* Operating Systems

You must have administrative rights to the operating system to install the drivers.

1. Install the adapter in the computer and turn on the computer.
2. Download the latest Dell Update Package (DUP) from Customer Support.
3. Run the DUP executable and click the **Install** button.
4. Follow the onscreen instructions.

Linux*

There are three methods for installing the Linux drivers:

- Install from Source Code
- Install from KMOD
- Install from KMP RPM

Please refer to the Linux section of this guide for more specific information.

Other Operating Systems

To install other drivers, visit the Customer Support web site: http://www.support.dell.com.

Installing the Adapter

Insert the PCI Express Adapter in the Server

NOTE: If you are replacing an existing adapter with a new adapter, you must re-install the driver.

1. Turn off the server and unplug the power cord, then remove the server's cover.

   CAUTION: Turn off and unplug the server before removing the server's cover. Failure to do so could endanger you and may damage the adapter or server.

2. Remove the cover bracket from an available PCI Express slot.

   NOTE: Some systems have physical x8 PCI Express slots that actually only support lower speeds. Please check your system manual to identify the slot.

3. Insert the adapter in an available, compatible PCI Express slot (refer to your card's specifications). Push the adapter into the slot until the adapter is firmly seated.
   You can install a smaller PCI Express adapter in a larger PCI Express slot.

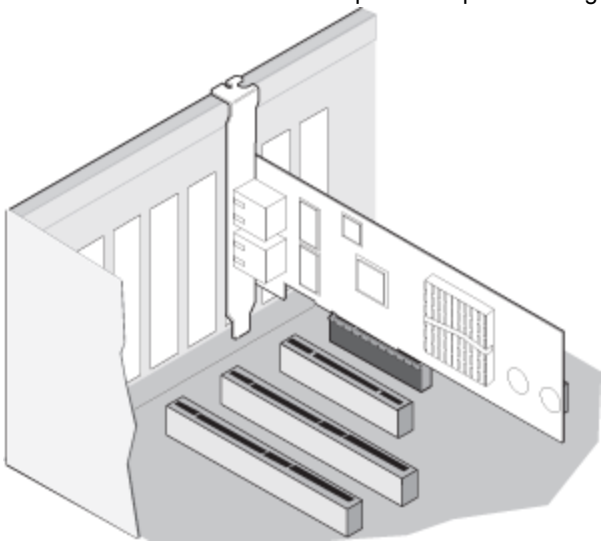

   CAUTION: Some PCI Express adapters may have a short connector, making them more fragile than PCI adapters. Excessive force could break the connector. Use caution when pressing the board in the slot.

4. Repeat steps 2 through 3 for each adapter you want to install.
5. Replace the server cover and plug in the power cord.
6. Turn the power on.

Connecting Network Cables

Connect the appropriate network cable, as described in the following sections.

Connect the UTP Network Cable

Insert the twisted pair, RJ-45 network cable as shown below.

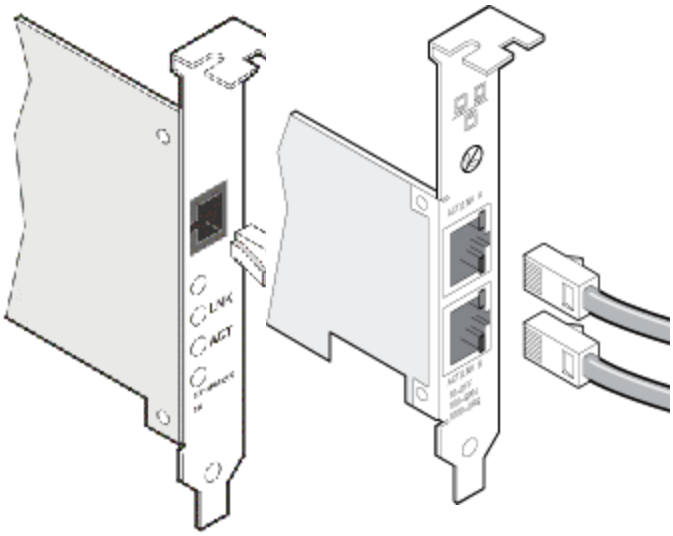

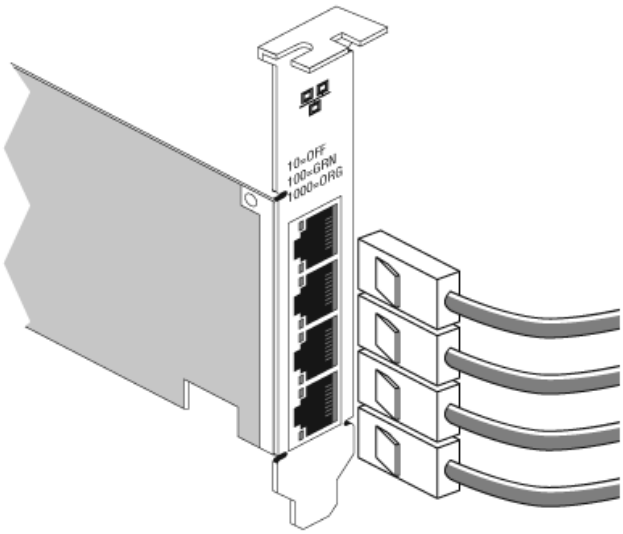

Single-port Adapter Dual-port Adapter Quad-port Adapter

Type of cabling to use:

- 10GBASE-T on Category 6, Category 6a, or Category 7 wiring, twisted 4-pair copper:
  - Length is 55 meters max for Category 6.
  - Length is 100 meters max for Category 6a.
  - Length is 100 meters max for Category 7.

  NOTE: For the Intel® 10 Gigabit AT Server Adapter, to ensure compliance with CISPR 24 and the EU's EN55024, this product should be used only with Category 6a shielded cables that are properly terminated according to the recommendations in EN50174-2.
- For 1000BASE-T or 100BASE-TX, use Category 5 or Category 5e wiring, twisted 4-pair copper:
  - Make sure you use Category 5 cabling that complies with the TIA-568 wiring specification. For more information on this specification, see the Telecommunications Industry Association's web site: www.tiaonline.org.
  - Length is 100 meters max.
  - Category 3 wiring supports only 10 Mbps.

  CAUTION: If using less than 4-pair cabling, you must manually configure the speed and duplex setting of the adapter and the link partner. In addition, with 2- and 3-pair cabling the adapter can only achieve speeds of up to 100Mbps.
- For 100BASE-TX, use Category 5 wiring.
- For 10Base-T, use Category 3 or 5 wiring.
- If you want to use this adapter in a residential environment (at any speed), use Category 5 wiring. If the cable runs between rooms or through walls and/or ceilings, it should be plenum-rated for fire safety.

In all cases:

- The adapter must be connected to a compatible link partner, preferably set to auto-negotiate speed and duplex for Intel gigabit adapters.
- Intel Gigabit and 10 Gigabit Server Adapters using copper connections automatically accommodate either MDI or MDI-X connections. The auto-MDI-X feature of Intel gigabit copper adapters allows you to directly connect two adapters without using a cross-over cable.

Connect the Fiber Optic Network Cable

CAUTION: The fiber optic ports contain a Class 1 laser device. When the ports are disconnected, always cover them with the provided plug. If an abnormal fault occurs, skin or eye damage may result if in close proximity to the exposed ports.

Remove and save the fiber optic connector cover. Insert a fiber optic cable into the ports on the network adapter bracket as shown below.

Most connectors and ports are keyed for proper orientation. If the cable you are using is not keyed, check to be sure the connector is oriented properly (transmit port connected to receive port on the link partner, and vice versa).

The adapter must be connected to a compatible link partner, such as an IEEE 802.3z-compliant gigabit switch, which is operating at the same laser wavelength as the adapter.

Conversion cables to other connector types (such as SC-to-LC) may be used if the cabling matches the optical specifications of the adapter, including length limitations.

Insert the fiber optic cable as shown below.

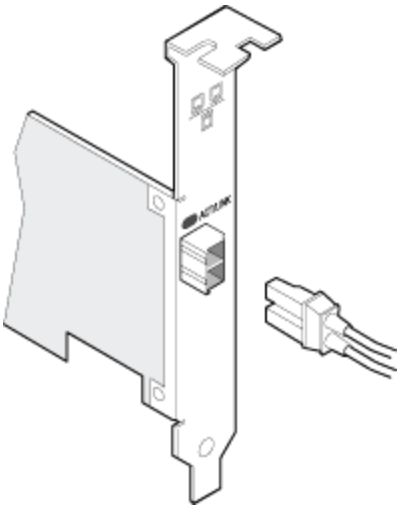

Connection requirements

- 40GBASE-SR4/MPO on 850 nanometer optical fiber:
  - Utilizing 50/125 micron OM3, length is 100 meters max.
  - Utilizing 50/125 micron OM4, length is 150 meters max.
- 10GBASE-SR/LC on 850 nanometer optical fiber:
  - Utilizing 50 micron multimode, length is 300 meters max.
  - Utilizing 62.5 micron multimode, length is 33 meters max.
- 1000BASE-SX/LC on 850 nanometer optical fiber:
  - Utilizing 50 micron multimode, length is 550 meters max.
  - Utilizing 62.5 micron multimode, length is 275 meters max.

SFP+ Devices with Pluggable Optics

Intel® Ethernet Server Adapters only support Intel optics and/or all passive and active limiting direct attach cables that comply with SFF-8431 v4.1 and SFF-8472 v10.4 specifications. When 82599-based SFP+ devices are connected back to back, they should be set to the same Speed setting using Intel PROSet for Windows or ethtool. Results may vary if you mix speed settings.

Supplier	Type	Part Numbers
Dell	Dual Rate 1G/10G SFP+ SR (bailed)	R8H2F, Y3KJN, 3G84K
Dell	TRIPLE RATE 1G/10G/40G QSFP+ SR (bailed) (1G not supported on XL710)	TCPM2, 27GG5, P8T4W

THIRD PARTY OPTIC MODULES AND CABLES REFERRED TO ABOVE ARE LISTED ONLY FOR THE PURPOSE OF HIGHLIGHTING THIRD PARTY SPECIFICATIONS AND POTENTIAL COMPATIBILITY, AND ARE NOT RECOMMENDATIONS OR ENDORSEMENT OR SPONSORSHIP OF ANY THIRD PARTY'S PRODUCT BY INTEL. INTEL IS NOT ENDORSING OR PROMOTING PRODUCTS MADE BY ANY THIRD PARTY AND THE THIRD PARTY REFERENCE IS PROVIDED ONLY TO SHARE INFORMATION REGARDING CERTAIN OPTIC MODULES AND CABLES WITH THE ABOVE SPECIFICATIONS. THERE MAY BE OTHER MANUFACTURERS OR SUPPLIERS, PRODUCING OR SUPPLYING OPTIC MODULES AND CABLES WITH SIMILAR OR MATCHING DESCRIPTIONS. CUSTOMERS MUST USE THEIR OWN DISCRETION AND DILIGENCE TO PURCHASE OPTIC MODULES AND CABLES FROM ANY THIRD PARTY OF THEIR CHOICE. CUSTOMERS ARE SOLELY RESPONSIBLE FOR ASSESSING THE SUITABILITY OF THE PRODUCT AND/OR DEVICES AND FOR THE SELECTION OF THE VENDOR FOR PURCHASING ANY PRODUCT. THE OPTIC MODULES AND CABLES REFERRED TO ABOVE ARE NOT WARRANTED OR SUPPORTED BY INTEL. INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF SUCH THIRD PARTY PRODUCTS OR SELECTION OF VENDOR BY CUSTOMERS.

Connect the Direct AttachTwinaxial Cable

Insert the twinaxial network cable as shown below.

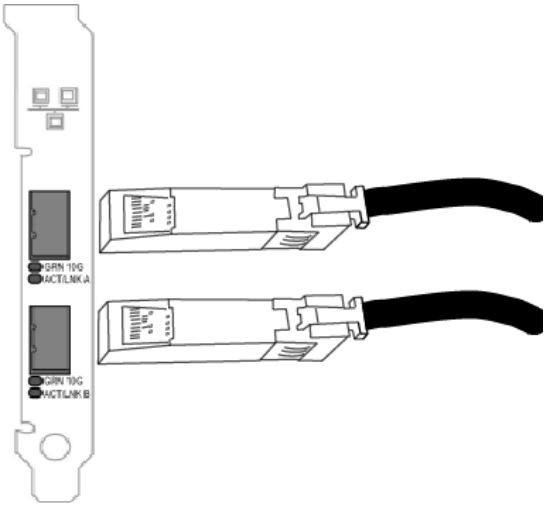

Type of cabling:

- 40 Gigabit Ethernet over SFP+ Direct Attached Cable (Twinaxial)
  - Length is 7 meters max.
- 10 Gigabit Ethernet over SFP+ Direct Attached Cable (Twinaxial)
  - Length is 10 meters max.

Install a Mezzanine Card in the Blade Server

See your server documentation for detailed instructions on how to install a Mezzanine card.

1. Turn off the blade server and pull it out of the chassis, then remove its cover.

   CAUTION: Failure to turn off the blade server could endanger you and may damage the card or server.

2. Lift the locking lever and insert the card in an available, compatible mezzanine card socket. Push the card into the socket until it is firmly seated.

   NOTE: A switch or pass-through module must be present on the same fabric as the card in the chassis to provide a physical connection. For example, if the mezzanine card is inserted in fabric B, a switch must also be present in fabric B of the chassis.

3. Repeat steps 2 for each card you want to install.
4. Lower the locking lever until it clicks into place over the card or cards.
5. Replace the blade server cover and put the blade back into the server chassis.
6. Turn the power on.

Install a Network Daughter Card in a Server

See your server documentation for detailed instructions on how to install a bNDC or rNDC.

1. Turn off the server and then remove its cover.

   CAUTION: Failure to turn off the server could endanger you and may damage the card or server.

2. Locate the Network Daughter Card connector in your server. See your server's documentation for details.
3. Press the Network Daughter Card into the connector.
4. Tighten the screws on the Network Daughter Card to secure it into place.
5. Replace the server's cover.

Setup

Installing Windows Network Drivers

Before you begin

To successfully install drivers or software, you must have administrative privileges on the computer.

Download the latest Dell Update Package from Customer Support.

Using the Dell Update Package (DUP)

The Dell Update Package (DUP) is an executable package that will update the network drivers on your system.

NOTES:

- If you are installing a driver in a computer with existing Intel adapters, be sure to update all the adapters and ports with the same driver and Intel® PROSet software. This ensures that all adapters will function correctly.
- If you have Fibre Channel over Ethernet (FCoE) boot enabled on any devices in the system, you will not be able to upgrade your drivers. You must disable FCoE boot before upgrading your Ethernet drivers.

Syntax

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe [/<option1>[=<value1>]] [/<option2>[=<value2>]]...

Command Line Option Descriptions

None	If you do not specify any command line options, the package will guide you through the installation.
/? or /h	Display the Update Package usage information.
/s	Suppress all graphical user interfaces of the Update Package.
/i	Do a fresh install of the drivers contained in the Update Package. **NOTE:** Requires /s option
/e=<path>	Extract the entire Update Package to the folder defined in <path>. **NOTE:** Requires /s option
/drivers=<path>	Extract only driver components of the Update Package to the folder defined in <path>. **NOTE:** Requires /s option
/driveronly	Install or Update only the driver components of the Update Package. **NOTE:** Requires /s option
/passthrough	(Advanced) Sends all text following the /passthrough option directly to the vendor install software of the Update Package. This mode suppresses any provided graphical user interfaces, but not necessarily those of the vendor software.
/capabilities	(Advanced) Returns a coded description of this Update Package's supported features. **NOTE:** Requires /s option
/l=<path>	Define a specific path for the Update Package log file. **NOTE:** This option can NOT be used in combination with /passthrough or /capabilities

/f	Override a soft dependency error returned from the Update Package. **NOTE:** Requires /s option, can NOT be used in combination with /passthrough or /capabilities

Examples

Update the system silently

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe /s

Fresh install silently

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe /s /i

Extract the update contents to the folder C:\mydir

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe /s /e=C:\mydir

Extract the driver components to the folder C:\mydir

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe /s /drivers=C:\mydir

Only install driver components

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe /s /driveronly

Change from the default log location to C:\my path with spaces\log.txt

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe /l="C:\my path with spaces\log.txt"

Force update to continue, even on "soft" qualification errors

Network_Driver_XXXXX_WN64_XX.X.X_A00.exe /s /f

Downgrading Drivers

You can use the /s and /f options to downgrade your drivers. For example, if you have the 17.0.0 drivers loaded and you want to downgrade to 16.5.0, type the following:

Network_Driver_XXXXX_WN64_16.5.0_A00.exe /s /f

Command Line Installation for Base Drivers and Intel® PROSet

Driver Installation

The driver install utility setup64.exe allows unattended installation of drivers from a command line.

These utilities can be used to install the base driver, intermediate driver, and all management applications for supported devices.

Setup64.exe Command Line Options

By setting the parameters in the command line, you can enable and disable management applications. If parameters are not specified, only existing components are updated.

Setup64.exe supports the following command line parameters:

Parameter	Definition
BD	Base Driver "0", do not install the base driver. "1", install the base driver (default).
ANS	Advanced Network Services "0", do not install ANS (default). If ANS is already installed, it will be uninstalled. "1", install ANS. The ANS property requires DMIX=1. **NOTE:** If the ANS parameter is set to ANS=1, both Intel PROSet and ANS will be installed.
DMIX	PROSet for Windows Device Manager "0", do not install Intel PROSet feature (default). If the Intel PROSet feature is already installed, it will be uninstalled. "1", install Intel PROSet feature. The DMIX property requires BD=1. **NOTE:** If DMIX=0, ANS will not be installed. If DMIX=0 and Intel PROSet, ANS, and FCoE are already installed, Intel PROSet, ANS, and FCoE will be uninstalled.
SNMP	Intel SNMP Agent "0", do not install SNMP (default). If SNMP is already installed, it will be uninstalled. "1", install SNMP. The SNMP property requires BD=1. **NOTE:** Although the default value for the SNMP parameter is 1 (install), the SNMP agent will only be installed if: • The Intel SNMP Agent is already installed. In this case, the SNMP agent will be updated. • The Windows SNMP service is installed. In this case, the SNMP window will pop up and you may cancel the installation if you do not want it installed.
FCOE	Fibre Channel over Ethernet "0", do not install FCoE (default). If FCoE is already installed, it will be uninstalled. "1", install FCoE. The FCoE property requires DMIX=1. **NOTE:** Even if FCOE=1 is passed, FCoE will not be installed if the operating system and installed adapters do not support FCoE.
ISCSI	iSCSI "0", do not install iSCSI (default). If iSCSI is already installed, it will be uninstalled. "1", install FCoE. The iSCSI property requires DMIX=1.
LOG	[log file name] LOG allows you to enter a file name for the installer log file. The default name is `C:\UmbInst.log`.
XML	[XML file name] XML allows you to enter a file name for the XML output file.
-a	Extract the components required for installing the base driver to `C:\Program Files\Intel\Drivers`. The directory where these files will be extracted to can be modified unless silent mode (/qn) is specified. If this parameter is specified, the installer will exit after the base driver is extracted. Any other parameters will be ignored.
-f	Force a downgrade of the components being installed. **NOTE:** If the installed version is newer than the current version, this parameter needs to be set.

<table>
<tr><th>Parameter</th><th>Definition</th></tr>
<tr><td>-v</td><td>Display the current install package version.</td></tr>
<tr><td>/q[r|n]</td><td>/q --- silent install options
<table>
<tr><td>r</td><td>Reduced GUI Install (only displays critical warning messages)</td></tr>
<tr><td>n</td><td>Silent install</td></tr>
</table></td></tr>
<tr><td>/l[i|w|e|a]</td><td>/l --- log file option for DMIX and SNMP installation. Following are log switches:
<table>
<tr><td>i</td><td>log status messages.</td></tr>
<tr><td>w</td><td>log non-fatal warnings.</td></tr>
<tr><td>e</td><td>log error messages.</td></tr>
<tr><td>a</td><td>log the start of all actions.</td></tr>
</table></td></tr>
<tr><td>-u</td><td>Uninstall the drivers.</td></tr>
</table>

NOTE: You must include a space between parameters.

Command line install examples

Assume that setup64.exe is in the root directory of the CD, D:\.

1. How to install the base driver:

   ```
   D:\Setup64.exe DMIX=0 ANS=0 SNMP=0
   ```

2. How to install the base driver using the LOG option:

   ```
   D:\Setup64.exe LOG=C:\installBD.log DMIX=0 ANS=0 SNMP=0
   ```

3. How to install Intel PROSet and ANS silently:

   ```
   D:\Setup64.exe DMIX=1 ANS=1 /qn
   ```

4. How to install Intel PROSet without ANS silently:

   ```
   D:\Setup64.exe DMIX=1 ANS=0 /qn
   ```

5. How to install components but deselect ANS:

   ```
   D:\Setup64.exe DMIX=1 ANS=0 /qn /liew C:\install.log
   ```

   The /liew log option provides a log file for the DMIX installation.

NOTE: To install teaming and VLAN support on a system that has adapter base drivers and Intel PROSet for Windows Device Manager installed, type the command line `D:\Setup64.exe ANS=1`.

Windows Server Core

In addition to the above method, on Windows Server Core, the base driver can also be installed using the Plug and Play Utility, PnPUtil.exe.

Installing Intel PROSet

Intel PROSet for Windows Device Manager is an advanced configuration utility that incorporates additional configuration and diagnostic features into the device manager. For information on installation and usage, see Using Intel® PROSet for Windows Device Manager.

NOTE: You must install Intel® PROSet for Windows Device Manager if you want to use Intel® ANS teams or VLANs.

Using the Adapter

Setting Speed and Duplex

Overview

The Link Speed and Duplex setting lets you choose how the adapter sends and receives data packets over the network.

In the default mode, an Intel network adapter using copper connections will attempt to auto-negotiate with its link partner to determine the best setting. If the adapter cannot establish link with the link partner using auto-negotiation, you may need to manually configure the adapter and link partner to the identical setting to establish link and pass packets. This should only be needed when attempting to link with an older switch that does not support auto-negotiation or one that has been forced to a specific speed or duplex mode.

Auto-negotiation is disabled by selecting a discrete speed and duplex mode in the adapter properties sheet.

NOTES:

- When an adapter is running in NPar mode, Speed settings are limited to the root partition of each port.
- Fiber-based adapters operate only in full duplex at their native speed.

The settings available when auto-negotiation is disabled are:

- **40 Gbps full duplex** (requires a full duplex link partner set to full duplex). The adapter can send and receive packets at the same time.
- **10 Gbps full duplex** (requires a full duplex link partner set to full duplex). The adapter can send and receive packets at the same time.
- **1 Gbps full duplex** (requires a full duplex link partner set to full duplex). The adapter can send and receive packets at the same time. You must set this mode manually (see below).
- **10 Mbps or 100 Mbps full duplex** (requires a link partner set to full duplex). The adapter can send and receive packets at the same time. You must set this mode manually (see below).
- **10 Mbps or 100 Mbps half duplex** (requires a link partner set to half duplex). The adapter performs one operation at a time; it either sends or receives. You must set this mode manually (see below).

Your link partner must match the setting you choose.

NOTES:

- Although some adapter property sheets (driver property settings) list 10 Mbps and 100 Mbps in full or half duplex as options, using those settings is not recommended.
- Only experienced network administrators should force speed and duplex manually.
- You cannot change the speed or duplex of Intel adapters that use fiber cabling.

Intel 10 Gigabit adapters that support 1 gigabit speed allow you to configure the speed setting. If this option is not present, your adapter only runs at its native speed.

Manually Configuring Duplex and Speed Settings

Configuration is specific to your operating system driver. To set a specific Link Speed and Duplex mode, refer to the section below that corresponds to your operating system.

CAUTION: The settings at the switch must always match the adapter settings. Adapter performance may suffer, or your adapter might not operate correctly if you configure the adapter differently from your switch.

Windows

The default setting is for auto-negotiation to be enabled. Only change this setting to match your link partner's speed and duplex setting if you are having trouble connecting.

1. In Windows Device Manager, double-click the adapter you want to configure.
2. On the **Link Speed** tab, select a speed and duplex option from the **Speed and Duplex** drop-down menu.
3. Click **OK**.

More specific instructions are available in the Intel PROSet help.

Linux

See Linux* Driver for the Intel® Gigabit Family of Adapters for information on configuring Speed and Duplex on Linux systems.

Testing the Adapter

Intel's diagnostic software lets you test the adapter to see if there are problems with the adapter hardware, the cabling, or the network connection.

Tests for Windows

Intel PROSet allows you to run four types of diagnostic tests.

- Connection Test: Verifies network connectivity by pinging the DHCP server, WINS server, and gateway.
- Cable Tests: Provide information about cable properties.

  NOTE: The Cable Test is not supported on all adapters. The Cable Test will only be available on adapters that support it.
- Hardware Tests: Determines if the adapter is functioning properly.

  NOTE: Hardware tests will fail if the adapter is configured for iSCSI Boot.

To access these tests, select the adapter in Windows Device Manager, click the **Link** tab, and click **Diagnostics**. A Diagnostics window displays tabs for each type of test. Click the appropriate tab and run the test.

The availability of these tests is hardware and operating system dependent.

DOS Diagnostics

Use the DIAGS test utility to test adapters under DOS.

Linux Diagnostics

The driver utilizes the ethtool interface for driver configuration and diagnostics, as well as displaying statistical information. ethtool version 1.6 or later is required for this functionality.

The latest release of ethtool can be found at: http://sourceforge.net/projects/gkernel.

NOTE: ethtool 1.6 only supports a limited set of ethtool options. Support for a more complete ethtool feature set can be enabled by upgrading ethtool to the latest version.

Responder Testing

The Intel adapter can send test messages to another Ethernet adapter on the same network. This testing is available in DOS via the diags.exe utility downloaded from Customer Support.

Adapter Teaming

ANS Teaming, a feature of the Intel® Advanced Network Services (Intel® ANS) component, lets you take advantage of multiple adapters in a system by grouping them together. Intel ANS can use features like fault tolerance and load balancing to increase throughput and reliability.

Teaming functionality is provided through the intermediate driver, Intel ANS. Teaming uses the intermediate driver to group physical adapters into a team that acts as a single virtual adapter. Intel ANS serves as a wrapper around one or more base drivers, providing an interface between the base driver and the network protocol stack. By doing so, the intermediate driver gains control over which packets are sent to which physical interface as well as control over other properties essential to teaming.

There are several teaming modes you can configure Intel ANS adapter teams to use.

Setting Up Adapter Teaming

Before you can set up adapter teaming in Windows*, you must install Intel® PROSet software. For more information on setting up teaming, see the information for your operating system.

Operating Systems Supported

The following links provide information on setting up teaming with your operating system:

- Windows

NOTE: To configure teams in Linux, use Channel Bonding, available in supported Linux kernels. For more information see the channel bonding documentation within the kernel source, located at Documentation/networking/bonding.txt.

Using Intel ANS Teams and VLANs inside a Guest Virtual Machine

Intel ANS Teams and VLANs are only supported in the following guest virtual machines

Host\Guest VM	Microsoft Windows Server 2008 R2 VM	Microsoft Windows Server 2012 R2 VM	Microsoft Windows Server 2012 R2 VM
Microsoft Windows Hyper-V	No Teams or VLANs	LBFO	LBFO
Linux Hypervisor (Xen or KVM)	ANS Teams and VLANs	LBFO ANS VLANs	LBFO ANS VLANs
VMware ESXi	ANS Teams and VLANs	LBFO ANS VLANs	LBFO ANS VLANs

Supported Adapters

Teaming options are supported on Intel server adapters. Selected adapters from other manufacturers are also supported. If you are using a Windows-based computer, adapters that appear in Intel PROSet may be included in a team.

NOTE: In order to use adapter teaming, you must have at least one Intel server adapter in your system. Furthermore, all adapters must be linked to the same switch or hub.

Conditions that may prevent you from teaming a device

During team creation or modification, the list of available team types or list of available devices may not include all team types or devices. This may be caused by any of several conditions, including:

- The operating system does not support the desired team type.
- The device does not support the desired team type or does not support teaming at all.
- The devices you want to team together use different driver versions.
- You are trying to team an Intel PRO/100 device with an Intel 10GbE device.
- You can add Intel® Active Management Technology (Intel® AMT) enabled devices to Adapter Fault Tolerance (AFT), Switch Fault Tolerance (SFT), and Adaptive Load Balancing (ALB) teams. All other team types are not supported. The Intel AMT enabled device must be designated as the primary adapter for the team.
- The device's MAC address is overridden by the Locally Administered Address advanced setting.
- Fibre Channel over Ethernet (FCoE) Boot has been enabled on the device.
- The device has “OS Controlled” selected on the Data Center tab.
- The device has a virtual NIC bound to it.
- The device is part of a Microsoft* Load Balancing and Failover (LBFO) team.

Configuration Notes

- Not all team types are available on all operating systems.
- Be sure to use the latest available drivers on all adapters.
- NDIS 6.2 introduced new RSS data structures and interfaces. Due to this, you cannot enable RSS on teams that contain a mix of adapters that support NDIS 6.2 RSS and adapters that do not.

- If you are using an Intel® 10GbE Server Adapter and an Intel® Gigabit adapter in the same machine, the driver for the Gigabit adapter must be updated at the same time as the Intel 10GbE adapter.
- If a team is bound to a Hyper-V virtual NIC, the Primary or Secondary adapter cannot be changed.
- Some advanced features, including hardware offloading, are automatically disabled when non-Intel adapters are team members to assure a common feature set.
- TOE (TCP Offload Engine) enabled devices cannot be added to an ANS team and will not appear in the list of available adapters.

  To enable teaming using Broadcom Advanced Control Suite 2:

  1. Load base drivers and Broadcom Advanced Control Suite 2 (always use the latest software releases from www.support.dell.com)
  2. Select the Broadcom device and go to the Advanced Tab
  3. Disable Receive Side Scaling
  4. Go to Resource Allocations and select TCP Offload Engine (TOE)
  5. Click on Configure and uncheck TCP Offload Engine (TOE) from the NDIS Configuration section

  To enable teaming using Broadcom Advanced Control Suite 3:

  1. Load base drivers and Broadcom Advanced Control Suite 3 (always use the latest software releases from www.support.dell.com)
  2. Select the Broadcom device and uncheck TOE from the Configurations Tab
  3. Click on Apply
  4. Choose NDIS entry of Broadcom device and disable Receive Side Scaling from Configurations Tab
  5. Click on Apply

- Spanning tree protocol (STP) should be disabled on switch ports connected to team adapters in order to prevent data loss when the primary adapter is returned to service (failback). Activation Delay is disabled by default. Alternatively, an activation delay may be configured on the adapters to prevent data loss when spanning tree is used. Set the Activation Delay on the advanced tab of team properties.
- Fibre Channel over Ethernet (FCoE)/Data Center Bridging (DCB) will be automatically disabled when an adapter is added to a team with non-FCoE/DCB capable adapters.
- ANS teaming of VF devices inside a Windows 2008 R2 guest running on an open source hypervisor is supported.
- An Intel® Active Management Technology (Intel AMT) enabled device can be added to Adapter Fault Tolerance (AFT), Switch Fault Tolerance (SFT), and Adaptive Load Balancing (ALB) teams. All other team types are not supported. The Intel AMT enabled device must be designated as the primary adapter for the team.
- Before creating a team, adding or removing team members, or changing advanced settings of a team member, make sure each team member has been configured similarly. Settings to check include VLANs and QoS Packet Tagging, Jumbo Packets, and the various offloads. These settings are available on the Advanced Settings tab. Pay particular attention when using different adapter models or adapter versions, as adapter capabilities vary.
- If team members implement Intel ANS features differently, failover and team functionality may be affected. To avoid team implementation issues:
  - Create teams that use similar adapter types and models.
  - Reload the team after adding an adapter or changing any Advanced features. One way to reload the team is to select a new preferred primary adapter. Although there will be a temporary loss of network connectivity as the team reconfigures, the team will maintain its network addressing schema.
- ANS allows you to create teams of one adapter. A one-adapter team will not take advantage of teaming features, but it will allow you to "hot-add" another adapter to the team without the loss of network connectivity that occurs when you create a new team.
- Before hot-adding a new member to a team, make sure that new member's link is down. When a port is added to a switch channel before the adapter is hot-added to the ANS team, disconnections will occur because the switch will start forwarding traffic to the port before the new team member is actually configured. The opposite, where the member is first hot-added to the ANS team and then added to the switch channel, is also problematic because ANS will forward traffic to the member before the port is added to the switch channel, and disconnection will occur.
- Intel 10 Gigabit Server Adapters can team with Intel Gigabit adapters and certain server-oriented models from other manufacturers. If you are using a Windows-based computer, adapters that appear in the Intel® PROSet teaming wizard may be included in a team.
- Network ports using OS2BMC should not be teamed with ports that have OS2BMC disabled.

- A reboot is required when any changes are made, such as modifying an advanced parameter setting of the base driver or creating a team or VLAN, on the network port that was used for a RIS install.
- Intel adapters that do not support Intel PROSet may still be included in a team. However, they are restricted in the same way non-Intel adapters are. See Multi-Vendor Teaming for more information.
- If you create a Multi-Vendor Team, you must manually verify that the RSS settings for all adapters in the team are the same.
- The table below provides a summary of support for Multi-Vendor Teaming.

Multi-vendor Teaming using Intel Teaming Driver (iANS/PROSet)		Teaming Mode Supported					Offload Support		Other Offload and RSS Support	
Intel	**Broadcom**	**AFT**	**SFT**	**ALB/RLB**	**SLA**	**LACP**	**LSO**	**CSO**	**TOE**	**RSS**
Intel PCI Express	Broadcom Device with TOE disabled	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Intel PCI Express	Broadcom Device with TOE enabled	No	No	No	No	No	No	No	No	No

Microsoft* Load Balancing and Failover (LBFO) teams

Intel ANS teaming and VLANs are not compatible with Microsoft's LBFO teams. Intel® PROSet will block a member of an LBFO team from being added to an Intel ANS team or VLAN. You should not add a port that is already part of an Intel ANS team or VLAN to an LBFO team, as this may cause system instability. If you use an ANS team member or VLAN in an LBFO team, perform the following procedure to restore your configuration:

1. Reboot the machine.
2. Remove LBFO team. Even though LBFO team creation failed, after a reboot Server Manager will report that LBFO is Enabled, and the LBFO interface is present in the 'NIC Teaming' GUI.
3. Remove the ANS teams and VLANs involved in the LBFO team and recreate them. This is an optional (all bindings are restored when the LBFO team is removed), but strongly recommended step.

NOTE: If you add an Intel AMT enabled port to an LBFO team, do not set the port to Standby in the LBFO team. If you set the port to Standby you may lose AMT functionality.

Teaming Modes

There are several teaming modes, and they can be grouped into these categories:

Fault Tolerance

Provides network connection redundancy by designating a primary controller and utilizing the remaining controllers as backups. Designed to ensure server availability to the network. When the user-specified primary adapter loses link, the iANS driver will "fail over" the traffic to the available secondary adapter. When the link of the primary adapter resumes, the iANS driver will "fail back" the traffic to the primary adapter. See Primary and Secondary Adapters for more information. The iANS driver uses link-based tolerance and probe packets to detect the network connection failures.

- Link-based tolerance - The teaming driver checks the link status of the local network interfaces belonging to the team members. Link-based tolerance provides fail over and fail back for the immediate link failures only.
- Probing - Probing is another mechanism used to maintain the status of the adapters in a fault tolerant team. Probe packets are sent to establish known, minimal traffic between adapters in a team. At each probe interval, each adapter in the team sends a probe packet to other adapters in the team. Probing provides fail over and fail back for immediate link failures as well as external network failures in the single network path of the probe packets between the team members.

Fault Tolerance teams include Adapter Fault Tolerance (AFT) and Switch Fault Tolerance (SFT).

Load Balancing

Provides transmission load balancing by dividing outgoing traffic among all the NICs, with the ability to shift traffic away from any NIC that goes out of service. Receive Load Balancing balances receive traffic.

Load Balancing teams include Adaptive Load Balancing (ALB) teams.

NOTE: If your network is configured to use a VLAN, make sure the load balancing team is configured to use the same VLAN.

Link Aggregation

Combines several physical channels into one logical channel. Link Aggregation is similar to Load Balancing.

Link Aggregation teams include Static Link Aggregation and IEEE 802.3ad: dynamic mode.

IMPORTANT

- **For optimal performance, you must disable the Spanning Tree Protocol (STP) on all the switches in the network when using AFT, ALB, or Static Link Aggregation teaming.**
- **When you create a team, a virtual adapter instance is created. In Windows, the virtual adapter appears in both the Device Manager and Network and Dial-up Connections. Each virtual adapter instance appears as "Intel Advanced Network Services Virtual Adapter." Do not attempt to modify (except to change protocol configuration) or remove these virtual adapter instances using Device Manager or Network and Dial-up Connections. Doing so might result in system anomalies.**
- **Before creating a team, adding or removing team members, or changing advanced settings of a team member, make sure each team member has been configured similarly. Settings to check include VLANs and QoS Packet Tagging, Jumbo Packets, and the various offloads. These settings are available in Intel PROSet's Advanced tab. Pay particular attention when using different adapter models or adapter versions, as adapter capabilities vary.**

If team members implement Advanced Features differently, failover and team functionality will be affected. To avoid team implementation issues:

- Use the latest available drivers on all adapters.
- Create teams that use similar adapter types and models.
- Reload the team after adding an adapter or changing any Advanced Features. One way to reload the team is to select a new preferred primary adapter. Although there will be a temporary loss of network connectivity as the team reconfigures, the team will maintain its network addressing schema.

Primary and Secondary Adapters

Teaming modes that do not require a switch with the same capabilities (AFT, SFT, ALB (with RLB)) use a primary adapter. In all of these modes except RLB, the primary is the only adapter that receives traffic. RLB is enabled by default on an ALB team.

If the primary adapter fails, another adapter will take over its duties. If you are using more than two adapters, and you want a specific adapter to take over if the primary fails, you must specify a secondary adapter. If an Intel AMT enabled device is part of a team, it must be designated as the primary adapter for the team.

There are two types of primary and secondary adapters:

- **Default primary adapter:** If you do not specify a preferred primary adapter, the software will choose an adapter of the highest capability (model and speed) to act as the default primary. If a failover occurs, another adapter becomes the primary. Once the problem with the original primary is resolved, the traffic will not automatically restore to the default (original) primary adapter in most modes. The adapter will, however, rejoin the team as a non-primary.
- **Preferred Primary/Secondary adapters:** You can specify a preferred adapter in Intel PROSet. Under normal conditions, the Primary adapter handles all traffic. The Secondary adapter will receive fallback traffic if the primary fails. If the Preferred Primary adapter fails, but is later restored to an active status, control is automatically switched back to the Preferred Primary adapter. Specifying primary and secondary adapters adds no benefit to SLA and IEEE 802.3ad dynamic teams, but doing so forces the team to use the primary adapter's MAC address.

To specify a preferred primary or secondary adapter in Windows

1. In the Team Properties dialog box's **Settings** tab, click **Modify Team**.
2. On the **Adapters** tab, select an adapter.
3. Click **Set Primary** or **Set Secondary**.

   NOTE: You must specify a primary adapter before you can specify a secondary adapter.

4. Click **OK**.

The adapter's preferred setting appears in the Priority column on Intel PROSet's **Team Configuration** tab. A "1" indicates a preferred primary adapter, and a "2" indicates a preferred secondary adapter.

Failover and Failback

When a link fails, either because of port or cable failure, team types that provide fault tolerance will continue to send and receive traffic. Failover is the initial transfer of traffic from the failed link to a good link. Failback occurs when the original adapter regains link. You can use the Activation Delay setting (located on the Advanced tab of the team's properties in Device Manager) to specify a how long the failover adapter waits before becoming active. If you don't want your team to failback when the original adapter gets link back, you can set the Allow Failback setting to disabled (located on the Advanced tab of the team's properties in Device Manager).

Adapter Fault Tolerance (AFT)

Adapter Fault Tolerance (AFT) provides automatic recovery from a link failure caused from a failure in an adapter, cable, switch, or port by redistributing the traffic load across a backup adapter.

Failures are detected automatically, and traffic rerouting takes place as soon as the failure is detected. The goal of AFT is to ensure that load redistribution takes place fast enough to prevent user sessions from being disconnected. AFT supports two to eight adapters per team. Only one active team member transmits and receives traffic. If this primary connection (cable, adapter, or port) fails, a secondary, or backup, adapter takes over. After a failover, if the connection to the user-specified primary adapter is restored, control passes automatically back to that primary adapter. For more information, see Primary and Secondary Adapters.

AFT is the default mode when a team is created. This mode does not provide load balancing.

NOTES

- AFT teaming requires that the switch not be set up for teaming and that spanning tree protocol is turned off for the switch port connected to the NIC or LOM on the server.
- All members of an AFT team must be connected to the same subnet.

Switch Fault Tolerance (SFT)

Switch Fault Tolerance (SFT) supports only two NICs in a team connected to two different switches. In SFT, one adapter is the primary adapter and one adapter is the secondary adapter. During normal operation, the secondary adapter is in standby mode. In standby, the adapter is inactive and waiting for failover to occur. It does not transmit or receive network traffic. If the primary adapter loses connectivity, the secondary adapter automatically takes over. When SFT teams are created, the Activation Delay is automatically set to 60 seconds.

In SFT mode, the two adapters creating the team can operate at different speeds.

NOTE: SFT teaming requires that the switch not be set up for teaming and that spanning tree protocol is turned on.

Configuration Monitoring

You can set up monitoring between an SFT team and up to five IP addresses. This allows you to detect link failure beyond the switch. You can ensure connection availability for several clients that you consider critical. If the connection between the primary adapter and all of the monitored IP addresses is lost, the team will failover to the secondary adapter.

Adaptive/Receive Load Balancing (ALB/RLB)

Adaptive Load Balancing (ALB) is a method for dynamic distribution of data traffic load among multiple physical channels. The purpose of ALB is to improve overall bandwidth and end station performance. In ALB, multiple links are provided from the server to the switch, and the intermediate driver running on the server performs the load balancing function. The ALB architecture utilizes knowledge of Layer 3 information to achieve optimum distribution of the server transmission load.

ALB is implemented by assigning one of the physical channels as Primary and all other physical channels as Secondary. Packets leaving the server can use any one of the physical channels, but incoming packets can only use the Primary Channel. With Receive Load Balancing (RLB) enabled, it balances IP receive traffic. The intermediate driver analyzes the send and transmit loading on each adapter and balances the rate across the adapters based on destination address. Adapter teams configured for ALB and RLB also provide the benefits of fault tolerance.

NOTES:

- ALB teaming requires that the switch not be set up for teaming and that spanning tree protocol is turned off for the switch port connected to the network adapter in the server.
- ALB does not balance traffic when protocols such as NetBEUI and IPX* are used.
- You may create an ALB team with mixed speed adapters. The load is balanced according to the adapter's capabilities and bandwidth of the channel.
- All members of ALB and RLB teams must be connected to the same subnet.

Virtual Machine Load Balancing

Virtual Machine Load Balancing (VMLB) provides transmit and receive traffic load balancing across Virtual Machines bound to the team interface, as well as fault tolerance in the event of switch port, cable, or adapter failure.

The driver analyzes the transmit and receive load on each member adapter and balances the traffic across member adapters. In a VMLB team, each Virtual Machine is associated with one team member for its TX and RX traffic.

If only one virtual NIC is bound to the team, or if Hyper-V is removed, then the VMLB team will act like an AFT team.

NOTES:

- VMLB does not load balance non-routed protocols such as NetBEUI and some IPX* traffic.
- VMLB supports from two to eight adapter ports per team.
- You can create a VMLB team with mixed speed adapters. The load is balanced according to the lowest common denominator of adapter capabilities and the bandwidth of the channel.
- An Intel AMT enabled adapter cannot be used in a VMLB team.

Static Link Aggregation

Static Link Aggregation (SLA) is very similar to ALB, taking several physical channels and combining them into a single logical channel.

This mode works with:

- Cisco EtherChannel capable switches with channeling mode set to "on"
- Intel switches capable of Link Aggregation
- Other switches capable of static 802.3ad

The Intel teaming driver supports Static Link Aggregation for:

- **Fast EtherChannel (FEC):** FEC is a trunking technology developed mainly to aggregate bandwidth between switches working in Fast Ethernet. Multiple switch ports can be grouped together to provide extra bandwidth. These aggregated ports together are called Fast EtherChannel. Switch software treats the grouped ports as a single logical port. An end node, such as a high-speed end server, can be connected to the switch using FEC. FEC link aggregation provides load balancing in a way which is very similar to ALB, including use of the same algorithm in the transmit flow. Receive load balancing is a function of the switch.

  The transmission speed will never exceed the adapter base speed to any single address (per specification). Teams must match the capability of the switch. Adapter teams configured for static Link Aggregation also provide the benefits of fault tolerance and load balancing. You do not need to set a primary adapter in this mode.

- **Gigabit EtherChannel (GEC):** GEC link aggregation is essentially the same as FEC link aggregation.

NOTES:

- All adapters in a Static Link Aggregation team must run at the same speed and must be connected to a Static Link Aggregation-capable switch. If the speed capabilities of adapters in a Static Link Aggregation team are different, the speed of the team is dependent on the switch.
- Static Link Aggregation teaming requires that the switch be set up for Static Link Aggregation teaming and that spanning tree protocol is turned off.
- An Intel AMT enabled adapter cannot be used in an SLA team.

IEEE 802.3ad: Dynamic Link Aggregation

IEEE 802.3ad is the IEEE standard. Teams can contain two to eight adapters. You must use 802.3ad switches (in dynamic mode, aggregation can go across switches). Adapter teams configured for IEEE 802.3ad also provide the benefits of fault tolerance and load balancing. Under 802.3ad, all protocols can be load balanced.

Dynamic mode supports multiple aggregators. Aggregators are formed by port speed connected to a switch. For example, a team can contain adapters running at 1 Gbps and 10 Gbps, but two aggregators will be formed, one for each speed. Also, if a team contains 1 Gbps ports connected to one switch, and a combination of 1 Gbps and 10 Gbps ports connected to a second switch, three aggregators would be formed. One containing all the ports connected to the first switch, one containing the 1 Gbps ports connected to the second switch, and the third containing the 10Gbps ports connected to the second switch.

NOTES:

- IEEE 802.3ad teaming requires that the switch be set up for IEEE 802.3ad (link aggregation) teaming and that spanning tree protocol is turned off.
- Once you choose an aggregator, it remains in force until all adapters in that aggregation team lose link.
- In some switches, copper and fiber adapters cannot belong to the same aggregator in an IEEE 802.3ad configuration. If there are copper and fiber adapters installed in a system, the switch might configure the copper adapters in one aggregator and the fiber-based adapters in another. If you experience this behavior, for best performance you should use either only copper-based or only fiber-based adapters in a system.
- An Intel AMT enabled adapter cannot be used in a DLA team.

Before you begin

- Verify that the switch fully supports the IEEE 802.3ad standard.
- Check your switch documentation for port dependencies. Some switches require pairing to start on a primary port.
- Check your speed and duplex settings to ensure the adapter and switch are running at full duplex, either forced or set to auto-negotiate. Both the adapter and the switch must have the same speed and duplex configuration. The full-duplex requirement is part of the IEEE 802.3ad specification: http://standards.ieee.org/. If needed, change your speed or duplex setting before you link the adapter to the switch. Although you can change speed and duplex settings after the team is created, Intel recommends you disconnect the cables until settings are in effect. In some cases, switches or servers might not appropriately recognize modified speed or duplex settings if settings are changed when there is an active link to the network.
- If you are configuring a VLAN, check your switch documentation for VLAN compatibility notes. Not all switches support simultaneous dynamic 802.3ad teams and VLANs. If you do choose to set up VLANs, configure teaming and VLAN settings on the adapter before you link the adapter to the switch. Setting up VLANs after the switch has created an active aggregator affects VLAN functionality.

Multi-Vendor Teaming

Multi-Vendor Teaming (MVT) allows teaming with a combination of Intel and non-Intel adapters.

If you are using a Windows-based computer, adapters that appear in the Intel PROSet teaming wizard can be included in a team.

MVT Design Considerations

- In order to activate MVT, you must have at least one Intel adapter or integrated connection in the team, which must be designated as the primary adapter.
- A multi-vendor team can be created for any team type.

- All members in an MVT must operate on a common feature set (lowest common denominator).
- For MVT teams, manually verify that the frame setting for the non-Intel adapter is the same as the frame settings for the Intel adapters.
- If a non-Intel adapter is added to a team, its RSS settings must match the Intel adapters in the team.

Virtual LANs

Overview

NOTE: Windows* users must install Intel® PROSet for Windows Device Manager and Advanced Networking Services in order to use VLANs.

The term VLAN (Virtual Local Area Network) refers to a collection of devices that communicate as if they were on the same physical LAN. Any set of ports (including all ports on the switch) can be considered a VLAN. LAN segments are not restricted by the hardware that physically connects them.

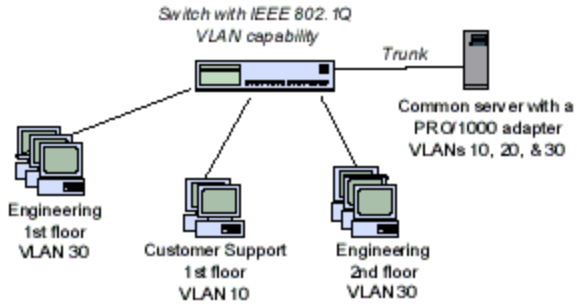

VLANs offer the ability to group computers together into logical workgroups. This can simplify network administration when connecting clients to servers that are geographically dispersed across the building, campus, or enterprise network.

Typically, VLANs consist of co-workers within the same department but in different locations, groups of users running the same network protocol, or a cross-functional team working on a joint project.

By using VLANs on your network, you can:

- Improve network performance
- Limit broadcast storms
- Improve LAN configuration updates (adds, moves, and changes)
- Minimize security problems
- Ease your management task

Other Considerations

- **Configuring SR-IOV for improved network security:** In a virtualized environment, on Intel® Server Adapters that support SR-IOV, the virtual function (VF) may be subject to malicious behavior. Software-generated frames are not expected and can throttle traffic between the host and the virtual switch, reducing performance. To resolve this issue, configure all SR-IOV enabled ports for VLAN tagging. This configuration allows unexpected, and potentially malicious, frames to be dropped.
- VLANs are not supported on Microsoft* Windows* 10. Any VLANs created with Release 20.1, 20.2 or 20.3 on a Windows 10 system will be corrupted and cannot be upgraded to version 20.4. The version 20.4 installer will remove existing VLANs.
- To set up IEEE VLAN membership (multiple VLANs), the adapter must be attached to a switch with IEEE 802.1Q VLAN capability.
- VLANs can co-exist with teaming (if the adapter supports both). If you do this, the team must be defined first, then you can set up your VLAN.
- You can set up only one untagged VLAN per adapter or team. You must have at least one tagged VLAN before you can set up an untagged VLAN.

IMPORTANT: When using IEEE 802.1Q VLANs, VLAN ID settings must match between the switch and those adapters using the VLANs.

NOTE: Intel ANS VLANs are not compatible with Microsoft's Load Balancing and Failover (LBFO) teams. Intel® PROSet will block a member of an LBFO team from being added to an Intel ANS VLAN. You should not add a port that is already part of an Intel ANS VLAN to an LBFO team, as this may cause system instability.

Configuring VLANs in Microsoft* Windows*

In Microsoft* Windows*, you must use Intel® PROSet to set up and configure VLANs. For more information, select Intel PROSet in the Table of Contents (left pane) of this window.

NOTES:

- VLANs are not supported on Microsoft* Windows* 10. Any VLANs created with Release 20.1, 20.2 or 20.3 on a Windows 10 system will be corrupted and cannot be upgraded to version 20.4. The version 20.4 installer will remove existing VLANs.
- If you change a setting under the Advanced tab for one VLAN, it changes the settings for all VLANS using that port.
- In most environments, a maximum of 64 VLANs per network port or team are supported by Intel PROSet.
- ANS VLANs are not supported on adapters and teams that have VMQ enabled. However, VLAN filtering with VMQ is supported via the Microsoft Hyper-V VLAN interface. For more information see Microsoft Hyper-V virtual NICs on teams and VLANs.
- You can have different VLAN tags on a child partition and its parent. Those settings are separate from one another, and can be different or the same. The only instance where the VLAN tag on the parent and child MUST be the same is if you want the parent and child partitions to be able to communicate with each other through that VLAN. For more information see Microsoft Hyper-V virtual NICs on teams and VLANs.

Advanced Features

NOTE: The options available on the Advanced tab are adapter and system dependent. Not all adapters will display all options.

Jumbo Frames

Jumbo Frames are Ethernet frames that are larger than 1518 bytes. You can use Jumbo Frames to reduce server CPU utilization and increase throughput. However, additional latency may be introduced.

NOTES:

- Jumbo Frames are supported at 1000 Mbps and higher. Using Jumbo Frames at 10 or 100 Mbps is not supported and may result in poor performance or loss of link.
- End-to-end network hardware must support this capability; otherwise, packets will be dropped.
- Intel adapters that support Jumbo Frames have a frame size limit of 9238 bytes, with a corresponding MTU size limit of 9216 bytes.

Jumbo Frames can be implemented simultaneously with VLANs and teaming.

NOTE: If an adapter that has Jumbo Frames enabled is added to an existing team that has Jumbo Frames disabled, the new adapter will operate with Jumbo Frames disabled. The new adapter's Jumbo Frames setting in Intel PROSet will not change, but it will assume the Jumbo Frames setting of the other adapters in the team.

To configure Jumbo Frames at the switch, consult your network administrator or switch user's guide.

Jumbo Frame Restrictions:

- Jumbo Frames are not supported in multi-vendor team configurations.
- Supported protocols are limited to IP (TCP, UDP).
- Jumbo Frames require compatible switch connections that forward Jumbo Frames. Contact your switch vendor for more information.
- The Jumbo Frame setting inside a virtual machine must be the same, or lower than, the setting on the physical port.
- When standard sized Ethernet frames (64 to 1518 bytes) are used, there is no benefit to configuring Jumbo Frames.
- The Jumbo Frames setting on the switch must be set to at least 8 bytes larger than the adapter setting for Microsoft* Windows* operating systems, and at least 22 bytes larger for all other operating systems.

For information on configuring Jumbo Frames in Windows, see the Intel PROSet for Windows Device Manager online help.

For information on configuring Jumbo Frames in Linux*, see the Linux Driver for the Intel Network Adapters.

Quality of Service

Quality of Service (QoS) allows the adapter to send and receive IEEE 802.3ac tagged frames. 802.3ac tagged frames include 802.1p priority-tagged frames and 802.1Q VLAN-tagged frames. In order to implement QoS, the adapter must be connected to a switch that supports and is configured for QoS. Priority-tagged frames allow programs that deal with real-time events to make the most efficient use of network bandwidth. High priority packets are processed before lower priority packets.

To implement QoS, the adapter must be connected to a switch that supports and is configured for 802.1p QoS.

QoS Tagging is enabled and disabled in the **Advanced** tab of Intel PROSet for Windows Device Manager.

Once QoS is enabled in Intel PROSet, you can specify levels of priority based on IEEE 802.1p/802.1Q frame tagging.

Data Center Bridging

Data Center Bridging (DCB) is a collection of standards-based extensions to classical Ethernet. It provides a lossless data center transport layer that enables the convergence of LANs and SANs onto a single unified fabric.

Furthermore, DCB is a configuration Quality of Service implementation in hardware. It uses the VLAN priority tag (802.1p) to filter traffic. That means that there are 8 different priorities that traffic can be filtered into. It also enables priority flow control (802.1Qbb) which can limit or eliminate the number of dropped packets during network stress. Bandwidth can be allocated to each of these priorities, which is enforced at the hardware level (802.1Qaz).

Adapter firmware implements LLDP and DCBX protocol agents as per 802.1AB and 802.1Qaz respectively. The firmware based DCBX agent runs in willing mode only and can accept settings from a DCBX capable peer. Software configuration of DCBX parameters via dcbtool/lldptool are not supported.

Saving and Restoring an Adapter's Configuration Settings

The Save and Restore Command Line Tool allows you to copy the current adapter and team settings into a standalone file (such as on a USB drive) as a backup measure. In the event of a hard drive failure, you can reinstate most of your former settings.

The system on which you restore network configuration settings must have the same configuration as the one on which the save was performed.

NOTES:

- Only adapter settings are saved (these include ANS teaming and VLANs). The adapter's driver is not saved.
- Restore using the script only once. Restoring multiple times may result in unstable configuration.
- The Restore operation requires the same OS as when the configuration was Saved.
- Intel® PROSet for Windows* Device Manager must be installed for the SaveRestore.ps1 script to run.
- For systems running a 64-bit OS, be sure to run the 64-bit version of Windows PowerShell, not the 32-bit (x86) version, when running the SaveRestore.ps1 script.

Command Line Syntax

SaveRestore.ps1 –Action save|restore [-ConfigPath] [-BDF]

SaveRestore.ps1 has the following command line options:

Option	Description
-Action	Required. Valid values: save | restore. The **save** option saves adapter and team settings that have been changed from the default settings. When you restore with the resulting file, any settings not contained in the file are assumed to be the default. The **restore** option restores the settings.

-ConfigPath	Optional. Specifies the path and filename of the main configuration save file. If not specified, it is the script path and default filename (saved_config.txt).
-BDF	Optional. Default configuration file names are saved_config.txt and Saved_StaticIP.txt. If you specify -BDF during a restore, the script attempts to restore the configuration based on the PCI Bus:Device:Function:Segment values of the saved configuration. If you removed, added, or moved a NIC to a different slot, this may result in the script applying the saved settings to a different device. **NOTES:** • If the restore system is not identical to the saved system, the script may not restore any settings when the -BDF option is specified. • Virtual Function devices do not support the -BDF option. • If you used Windows to set NPar minimum and maximum bandwidth percentages, you must specify /bdf during save and restore to keep those settings.

Examples

Save Example

To save the adapter settings to a file on a removable media device, do the following.

1. Open a Windows PowerShell Prompt.
2. Navigate to the directory where SaveRestore.ps1 is located (generally c:\Program Files\Intel\DMIX).
3. Type the following:
   ```
   SaveRestore.ps1 –Action Save –ConfigPath e:\settings.txt
   ```

Restore Example

To restore the adapter settings from a file on removable media, do the following:

1. Open a Windows PowerShell Prompt.
2. Navigate to the directory where SaveRestore.ps1 is located (generally c:\Program Files\Intel\DMIX).
3. Type the following:
   ```
   SaveRestore.ps1 –Action Restore –ConfigPath e:\settings.txt
   ```

NIC Partitioning

Network Interface Card (NIC) Partitioning (NPar) allows network administrators to create multiple partitions for each physical port on a network adapter card, and to set different bandwidth allocations on each partition. To the network and operating system, each partition appears as a separate physical port on the adapter. This facilitates the reduction of switch port count and cabling complexity while maintaining network segmentation and isolation. In addition, flexible bandwidth allocation per partition allows for efficient use of the link.

NPar is available in Linux and in Windows Server and Windows Server Core versions starting with 2008 R2.

NPar supports a maximum of 8 partitions on the following adapters.

- Intel® Ethernet 10G 4P X710/I350 rNDC
- Intel® Ethernet 10G 4P X710-k bNDC
- Intel® Ethernet 10G 2P X710-k bNDC
- Intel® Ethernet 10G X710-k bNDC
- Intel® Converged Network Adapter X710

NOTES:

- Adapters support NPar in NIC (LAN) mode only.
- The following are supported on the first partition of each port only:
  - PXE Boot
  - iSCSIboot
  - Speed and Duplex settings
  - Flow Control
  - Power Management settings
  - SR-IOV
  - NVGRE processing
- Resource limits in Microsoft Windows may affect the number of ports that are displayed. If you have several adapters installed in a system, and enable NPar or NParEP on the adapters, Windows Device Manager may not display all of the ports.
- On an X710 device in a Microsoft Windows Server 2008 R2 system with NParEP enabled, the driver will load on the first 8 physical functions only. The second set of 8 physical functions will have Code 10 yellow bang errors in Windows Device Manager. This is a limitation of the operating system.

NParEP Mode

NParEP Mode is a combination of NPar and PCIe ARI, and increases the maximum number of partitions on an adapter to 16 per NIC.

NParEP Platform Support

		PCI Express Slot									
Dell Platform	**Rack NDC Slot**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
C4130		yes	yes								
R230		no	no								
R330		no	no								
R430		yes	yes								
R530		yes	yes	yes	no	no					
R530XD		yes	yes	no							
R630	yes	yes	yes	yes							
R730	yes	yes	yes	yes	yes	yes	yes	yes			

		PCI Express Slot									
Dell Platform	**Rack NDC Slot**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
R730XD	yes	yes	yes	yes	yes	yes	yes				
R830	yes	yes	yes	yes	yes	yes	yes				
R930	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
T130		no	no	no	no						
T330		no	no	no	yes						
T430		no	no	yes	yes	yes	yes				
T630		yes	no	yes	yes	yes	yes	yes			

		Mezzanine Slot	
Dell Platform	**Blade NDC Slot**	**B**	**C**
FC430			
FC630	yes		
FC830	yes		
M630	yes		
M630 for VRTX	yes		
M830	yes		
M830 for VRTX	yes		

Supported platforms or slots are indicated by "yes." Unsupported are indicated by "no". Not applicable are indicated by blank cells.

Configuring NPar Mode

Configuring NPar from the Boot Manager

When you boot the system, press the **F2** key to enter the **System Setup** menu. Select **Device Settings** from the list under **System Setup Main Menu**, then select your adapter from the list to get to the Device Configuration menu. Select **Device Level Configuration** in the list under **Main Configuration Page**. This brings up the Virtualization settings under **Device Level Configuration**.

There are four options in the **Virtualization Mode** drop down list.

- None: the adapter operates normally
- NPar: allows up to 8 partitions on the adapter. If you select NPar Virtualization Mode, you will then be presented with the option to enable NParEP Mode, which will extend the number of partitions per adapter to a total of 16 by pairing NPar with PCIe ARI.

  NOTES:
  - When an adapter is running in NPar Mode, it is limited to 8 partitions total. A two-port adapter will have four partitions per port. A four-port adapter will have two partitions per port.
  - NParEP Mode can only be enabled when NPar Mode has been enabled.
  - When an adapter is running in NParEP Mode, it is limited to 16 partitions total. A two-port adapter will have eight partitions per port. A four port adapter will have four partitions per port.
- SR-IOV: activates SR-IOV on the port
- NPar+SR-IOV: Allows up to 8 partitions (physical functions) for the adapter and activates SR-IOV.

  NOTES:
  - SR-IOV is limited to the root partition of each port.
  - When an adapter is running in NPar mode, virtualization (SR-IOV) settings apply to all ports on the adapter, and to all partitions on each port. Changes made the virtualization settings on one port are applied to all ports on the adapter.

When you have completed your selection, click the **Back** button, and you will return to the **Main Configuration Page**. Click the new item, titled **NIC Partitioning Configuration**, in the configuration list to go to the NIC Partitioning Configuration page, where you will see a list of the NPar (or NParEP) partitions on your adapter.

The Global Bandwidth Allocation page lets you specify the minimum and maximum guaranteed bandwidth allocation for each partition on a port. Minimum TX Bandwidth is the guaranteed minimum data transmission bandwidth, as a percentage of the full physical port link speed, that the partition will receive. The bandwidth the partition is awarded will never fall below the level you specify here. The valid range of values is:

1 to ((100 minus # of partitions on the physical port) plus 1)

For example, if a physical port has 4 partitions, the range would be:

1 to ((100 - 4) + 1 = 97)

The Maximum Bandwidth percentage represents the maximum transmit bandwidth allocated to the partition as a percentage of the full physical port link speed.The accepted range of values is 0-100. The value here can be used as a limiter, should you chose that any one particular partition not be able to consume 100% of a port's bandwidth should it be available. The sum of all the values for Maximum Bandwidth is not restricted, because no more than 100% of a port's bandwidth can ever be used.

Click the **Back** button when you have finished making your bandwidth allocation settings to return to the NIC Partitioning Configuration page. From there you may click on one of the **Partition*n*Configuration** list items under **Global Bandwidth Allocation**. This will bring up a partition configuration information page on a particular port. You can see the NIC Mode, PCI;Device ID, PCI Address, MAC Address, and the virtual MAC Address (if applicable) for all the partitions on any given port by clicking through the items in the Partition Configuration list.

When you have completed the configuration of all the partitions on a port, back out to the Main Configuration Page, click the **Finish** button, then click the **OK** button in the Success (Saving Changes) dialog.

Repeat the partition configuration process for all the ports on your adapter.

NOTE: Once NPar has been enabled on one of the partition on a port, it will appear enabled for all subsequent partitions on that port. If that first setting of NPar included enabling NParEP mode, NParEP Mode will appear enabled on all subsequent partitions on that port as well.

When you have configured all the partitions on all the ports on all the adapters in your server, back out to the System Setup Main Menu, and click the **Finish** button. Then click **Yes** to exit the System Setup Menu and to reboot the system in order to apply your changes.

Once the system has completed the boot process, NPar will remain enabled until you explicitly disable it by turning off the option during a subsequent boot sequence.

Configuring NPar in Microsoft Windows*

You can configure an adapter port partition in Windows just like any adapter port. Run Device Manager, select the and open the partition's properties sheets to configure options.

Enabling NPar

NPar is enabled and disabled from the **Advanced** tab in the Device Manager property sheet.

Boot Options

On the Boot Options tab, you will be advised that the device is in NPar mode and that legacy preboot protocol settings can only be configured on the root partition. Clicking the **Properties** button will launch the property sheet for the root partition on the adapter.

Power Management Settings

Power Management settings are allowed only on the first partition of each physical port. If you select the **Power Management** tab in the Device Manager property sheets while any partition other than the first partition is selected, you will be presented with text in the Power Management dialog stating that Power Management settings cannot be configured on the current connection. Clicking the **Properties** button will launch the property sheet for the root partition on the adapter.

NOTE: Boot options and Power Management settings are only available on the root partition of each physical port.

Flow Control

You can change the Flow Control settings for any partition on a given port. However, when a change is made to the Flow Control settings of a partition associated with a port on an adapter operating in NPar mode, the new value will be applied to all partitions on that particular port.

Flow control is reached by selecting the Intel PROSet **Advanced** tab, then selecting the **Properties** button, and then selecting **Flow Control** from the list of options in the **Settings** list of the dialog that is displayed.

Identifying Port Associations

The Hardware Information dialog in the Intel PROSet property sheets facilitates the identification of the physical port associated with a particular partition. There is an **Identify Adapter** button in the **Link Speed** tab, and clicking that button will cause the ACK/Link light on the port associated with the active partition to blink.

Partition Bandwidth Configuration

The Bandwidth Configuration dialog provides an indication of the port for which settings are currently being made, above a list of the partitions on that port and their current bandwidth allocations (Min%, Max%). Partition Bandwidth Configuration is reached by clicking the **Bandwidth Configuration** button on the **Link Speed** tab in the Intel PROSet properties sheets.

The bandwidth allocated to each partition on the port will never drop below the value set under Min%. For all the partitions on the same physical port, the min bandwidth percentage for all of the partitions must be set to zero, or the sum of all of the minimum bandwidth percentages on each partition must equal 100, where the range of min bandwidth percentages is between 1 and (100-n)%, where *n* is the number of partitions for a particular port. For example, on a port with four defined partitions:

P1=0 P2=0 P3=0 P4=0 Valid	P1=10 P2=20 P3=30 P4=40 Valid	P1=20 P2=80 P3=0 P4=0 NOT Valid

Valid values for Max% are the value of that partition's "Min%" through "100". For example, if the Min% value for Partition 1 is 50%, the range for that partition's Max% is "50"-"100". If you cause any one partition's Max% value to exceed 100% by incrementing the value with the spinner, an error is displayed and that Max% is decremented to 100%. The *sum* of the Max% values for all partitions on a particular port has no limit.

To change the value for Min% or Max%, select a partition in the displayed list, then use the up or down arrows under "Selected Partition Bandwidth Percentages".

Speed and Duplex Settings

The Speed and Duplex setting for a particular port may be changed from any partition associated with that port. However, because all partitions on a particular port on an adapter operating in NPar mode share the same module that is plugged into the port, changing the Speed and Duplex setting will result in the new value being set across all partitions on that same physical port.

Changing the Speed and Duplex setting for a port on an adapter running in NPar mode will cause the driver for each partition associated with that port to be reloaded. This may result in a momentary loss of link.

Online Diagnostics

Online tests can be performed while in NPar mode without the adapter losing link. The following diagnostics tests are available for all partitions for a particular port while an adapter is running in NPar mode:

- EEPROM
- Register
- NVM Integrity
- Connection

Offline Diagnostics

Offline diagnostics are not supported while an adapter is running in NPar mode. Loopback tests and Cable Offline tests are not allowed in NPar mode.

NPar Teaming Rules

No two ANS team member partitions can exist that are bound to the same physical port. If you attempt to add a partition to an existing team via the Teaming tab in Intel PROSet properties sheets for an adapter running in NPar mode, a check will be made as to whether the partition to be added is bound to the same physical port as any pre-existing team member.

When an adapter is added to a team, setting changes for the adapter and the team may cause momentary loss of connectivity.

Virtualization

Settings for virtualization (Virtual Machine Queues and SR-IOV) are reached through the Intel PROSet Properties sheets by selecting the Advanced tab, then selecting "Virtualization" from the Settings list.

When an adapter is operating in NPar mode, only the first partition of each physical port may be configured with the virtualization settings.

NOTE: Microsoft* Hyper-V* must be installed on the system in order for virtualization settings to be available. Without Hyper-V* being installed, the Virtualization tab in PROSet will not appear.

Configuring NPar in Linux*

See the Linux driver README file for details.

Exiting NPar Mode

NPar mode is disabled in the System Setup menu during a reboot.

Reboot the system, and press the **F2** key to enter the **System Setup** menu. Select **Device Settings** from the list under **System Setup Main Menu**, then select your adapter from the list to get to the Device Configuration menu. Select **Device Level Configuration** in the list under **Main Configuration Page**. This brings up the Virtualization settings under **Device Level Configuration**.

In the Virtualization Mode list, select "None". Then click the **Back** button, which returns you to the Main Configuration Page. There, click the **Finish** button to save your change and reboot the system. When the system completes the reboot, NPar will no longer be active.

NOTE: When NPar has been disabled and the system completes the reboot, any other virtualization-related settings, such as NParEP or SR-IOV will also be disabled.

Network Virtualization using Generic Routing Encapsulation (NVGRE)

Network Virtualization using Generic Routing Encapsulation (NVGRE) increases the efficient routing of network traffic within a virtualized or cloud environment. Some Intel® Ethernet Network devices perform Network Virtualization using Generic Routing Encapsulation (NVGRE) processing, offloading it from the operating system. This reduces CPU utilization.

NOTE: When a port is in NPar mode, NVGRE (the Encapsulated Task Offload setting) is available only on the first partition on the port.

Remote Wake-Up

Remote wake-up can wake your server from a low power or powered off state. If Wake On LAN is enabled, when your system is powered down, the network interface draws standby power and listens for specially designed packet. If it receives such a packet it will wake your server.

Advanced Configuration and Power Interface (ACPI)

ACPI supports a variety of power states. Each state represents a different level of power, from fully powered up to completely powered down, with partial levels of power in each intermediate state.

ACPI Power States

Power State	Description
S0	On and fully operational
S1	System is in low-power mode (sleep mode). The CPU clock is stopped, but RAM is powered on and being refreshed.
S2	Similar to S1, but power is removed from the CPU.
S3	Suspend to RAM (standby mode). Most components are shut down. RAM remains operational.
S4	Suspend to disk (hibernate mode). The memory contents are swapped to the disk drive and then reloaded into RAM when the system is awakened.
S5	Power off

Supported Adapters

All adapter in this release support Wake On LAN. The following adapters only support it on Port A:

- Intel® Gigabit 2P I350-t Adapter (port A only)
- Intel® Gigabit 4P I350-t Adapter (port A only)

NOTES:

- Not all systems support every wake setting. There may be BIOS or operating system settings that need to be enabled for your system to wake up. In particular, this is true for Wake from S5 (also referred to as Wake from power off).
- When a port is operating in NPar mode, WoL is available only on the first partition of each port.

Enabling Wake From Power Off

If you want to wake your system from a power off state, you must enable it from the System Setup.

1. Go to System Setup.
2. Choose a port and go to configuration.
3. Specify Wake on LAN.

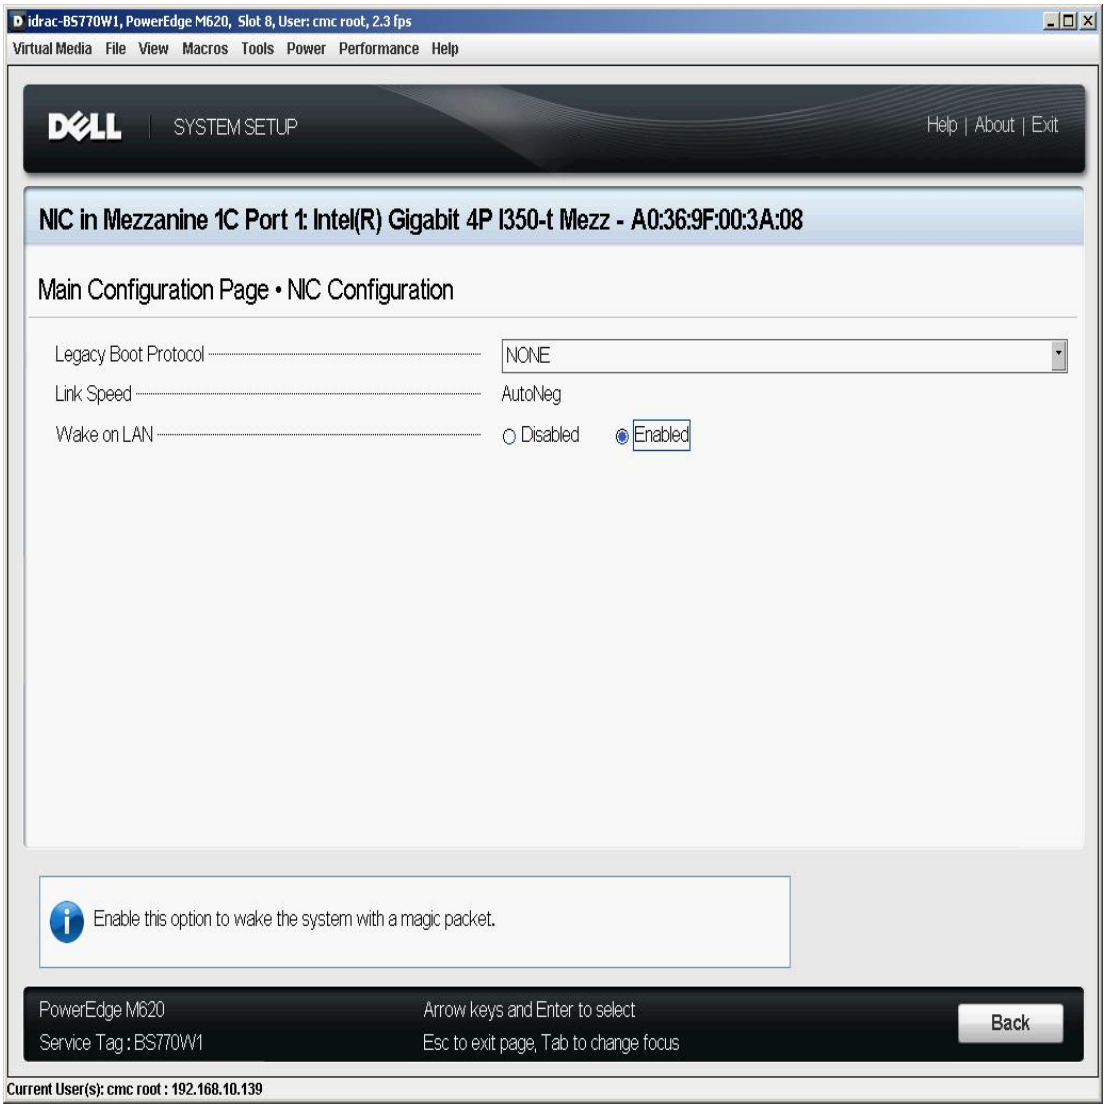

Wake-Up Address Patterns

Remote wake-up can be initiated by a variety of user selectable packet types and is not limited to the Magic Packet format. For more information about supported packet types, see the operating system settings section.

The wake-up capability of Intel adapters is based on patterns sent by the OS. You can configure the driver to the following settings using Intel PROSet for Windows. For Linux*, WoL is provided through the ethtool* utility. For more information on ethtool, see the following Web site: http://sourceforge.net/projects/gkernel.

- Wake on Directed Packet - accepts only patterns containing the adapter's Ethernet address in the Ethernet header or containing the IP address, assigned to the adapter, in the IP header.
- Wake on Magic Packet - accept only patterns containing 16 consecutive repetitions of the adapter's MAC address.
- Wake on Directed Packet and Wake on Magic Packet - accepts the patterns of both directed packets and magic packets.

Choosing "Wake on directed packet" will also allow the adapter to accept patterns of the Address Resolution Protocol (ARP) querying the IP address assigned to the adapter. If multiple IP addresses are assigned to an adapter, the operating system may request to wake up on ARP patterns querying any of the assigned addresses. However, the adapter will only awaken in response to ARP packets querying the first IP address in the list, usually the first address assigned to the adapter.

Physical Installation Issues

Slot

Some motherboards will only support remote wake-up (or remote wake-up from S5 state) in a particular slot. See the documentation that came with your system for details on remote wake-up support.

Power

Newer Intel PRO adapters are 3.3 volt and some are 12 volt. They are keyed to fit either type of slot.

The 3.3 volt standby supply must be capable of supplying at least 0.2 amps for each Intel PRO adapter installed. Turning off the remote wake-up capability on the adapter using the BootUtil utility reduces the power draw to around 50 milliamps (.05 amps) per adapter.

Operating System Settings

Microsoft Windows Products

Windows Server is ACPI-capable. These operating systems do not support remote wake-up from a powered off state (S5), only from standby. When shutting down the system, they shut down ACPI devices including the Intel PRO adapters. This disarms the adapters' remote wake-up capability. However, in some ACPI-capable computers, the BIOS may have a setting that allows you to override the OS and wake from an S5 state anyway. If there is no support for wake from S5 state in your BIOS settings, you are limited to wake from standby when using these operating systems in ACPI computers.

The **Power Management** tab in Intel PROSet includes a setting called Wake on Magic Packet from power off state for some adapters. To explicitly allow wake-up with a Magic Packet from shutdown under APM power management mode, check this box to enable this setting. See Intel PROSet help for more details.

In ACPI-capable versions of Windows, the Intel PROSet advanced settings include a setting called Wake on Settings. This setting controls the type of packets that wake the system from standby. See Intel PROSet help for more details.

If you do not have Intel PROSet installed you will need to do the following:

1. Open the Device Manager, then navigate to the **Power Management** tab, check "**Allow this device to bring the computer out of standby.**"

2. On the **Advanced** tab, set the "**Wake on Magic packet**" option to Enabled.

In order to wake from S5 without Intel PROSET, on the **Advanced tab**, set "**Enable PME**" to Enabled.

Other Operating Systems

Remote Wake-Up is also supported in Linux.

Optimizing Performance

You can configure Intel network adapter advanced settings to help optimize server performance.

The examples below provide guidance for three server usage models:

- Optimized for quick response and low latency – useful for video, audio, and High Performance Computing Cluster (HPCC) servers
- Optimized for throughput – useful for data backup/retrieval and file servers
- Optimized for CPU utilization – useful for application, web, mail, and database servers

NOTES:

- The recommendations below are guidelines and should be treated as such. Additional factors such as installed applications, bus type, network topology, and operating system also affect system performance.

- These adjustments should be performed by a highly skilled network administrator. They are not guaranteed to improve performance. Not all settings shown here may be available through your BIOS, operating system or network driver configuration. Linux users, see the README file in the Linux driver package for Linux-specific performance enhancement details.
- When using performance test software, refer to the documentation of the application for optimal results.

1. Install the adapter in a PCI Express bus slot.
2. Use the proper fiber cabling for the adapter you have.
3. Enable Jumbo Packets, if your other network components can also be configured for it.
4. Increase the number of TCP and Socket resources from the default value. For Windows based systems, we have not identified system parameters other than the TCP Window Size which significantly impact performance.

5. Increase the allocation size of Driver Resources (transmit/receive buffers). However, most TCP traffic patterns work best with the transmit buffer set to its default value, and the receive buffer set to its minimum value.

For specific information on any advanced settings, see Advanced Settings for Windows* Drivers or Linux* Driver for the Intel® Network Server Adapters.

Optimized for quick response and low latency

- Minimize or disable Interrupt Moderation Rate.
- Disable Offload TCP Segmentation.
- Disable Jumbo Packets.
- Increase Transmit Descriptors.
- Increase Receive Descriptors.
- Increase RSS Queues.

Optimized for throughput

- Enable Jumbo Packets.
- Increase Transmit Descriptors.
- Increase Receive Descriptors.
- On systems that support NUMA, set the Preferred NUMA Node on each adapter to achieve better scaling across NUMA nodes.

Optimized for CPU utilization

- Maximize Interrupt Moderation Rate.
- Keep the default setting for the number of Receive Descriptors; avoid setting large numbers of Receive Descriptors.
- Decrease RSS Queues.
- In Hyper-V environments, decrease the Max number of RSS CPUs.

Windows Drivers

Installing Windows* Drivers

Installing the Drivers

The drivers can be installed using the Found New Hardware wizard.

Installing Drivers on Windows Server Using the Found New Hardware Wizard

NOTES:

- When Windows Server detects a new adapter, it attempts to find an acceptable Windows driver already installed on the computer. If the operating system finds a driver, it installs this driver without any user intervention. However, this Windows driver may not be the most current one and may provide only basic functionality. Update the driver to make sure you have access to all the base driver's features.
- The Roll Back Driver feature of Windows Server (available on the Adapter Properties dialog's **Driver** tab) will not work correctly if an adapter team or Intel PROSet are present on the system.
  Before you use the Roll Back Driver feature, use Intel PROSet to remove any teams, then remove Intel PROSet using **Programs and Features** from the Control Panel of Windows.
- Using Microsoft Windows Update to upgrade or downgrade your Ethernet network drivers is not supported. Please download the latest driver package from http://www.dell.com/support.

1. Install the adapter in the computer and turn on the computer.
2. When Windows discovers the new adapter, the Found New Hardware Wizard starts.
3. Extract the Dell Driver Update Package to a specified path.
4. Open a DOS command box and go to the specified path.
5. Type "setup -a" at the command prompt to extract the drivers.
6. Type in the directory path where you want the files saved. The default path is c:\Program Files\Intel\Drivers.
7. The Wizard Welcome screen asks whether you want to connect to Windows Update to search for software.
   Click **No, not this time**.
   Click **Next**.
8. Click **Install from a list or specific location**, then click **Next**.
9. On the next screen, type in the directory path where you saved the driver files and click **Next**.
10. Windows searches for a driver. When the search is complete, a message indicates a driver was found.
11. Click **Next**.
    The necessary files are copied to your computer. The wizard displays a Completed message.
12. Click **Finish**.

If Windows does not detect the adapter, see Troubleshooting.

Installing Drivers Using the Windows Command Line

You can also use the Windows command line to install the drivers. The driver install utility (setup64.exe) allows unattended install of the drivers.

For complete information, see Command Line Installation for Base Drivers and Intel® PROSet.

Installing Additional Adapters

When you use the Found New Hardware Wizard to install drivers, Windows installs the driver for the first adapter and then automatically installs drivers for additional adapters.

There are no special instructions for installing drivers of non-Intel adapters (e.g., for multi-vendor teaming). Follow the instructions that came with that adapter.

Updating the Drivers

NOTE: If you update the adapter driver and are using Intel PROSet, you should also update Intel PROSet. To update the application, double-click setup64.exe and make sure the option for Intel® PROSet for Windows Device Manager is checked.

Drivers can be updated using the Update Device Driver Wizard.

Updating Windows Server Using the Device Manager

1. Extract the Dell Driver Update Package to a specified path.
2. From the Control Panel, double-click the **System** icon and click **Device Manager**.
3. Double-click **Network Adapters** and right-click on the Intel adapter listing to display its menu.
4. Click the **Update Driver...** menu option. The Update Driver Software page appears.
5. Select **Browse my computer for driver software**.
6. Type in the directory path to the specified drive or browse to the location.
7. Select **Next**.
8. After the system finds and installs the file, click **Close**.

Removing the Drivers

You should uninstall the Intel driver if you are permanently removing all Intel adapters, or if you need to perform a clean installation of new drivers. This procedure removes the driver for all Intel adapters that use it as well as Intel PROSet and Advanced Networking Services.

WARNING: Removing an adapter driver results in a disruption of all network traffic through that adapter.

NOTE: Before you remove the driver, make sure the adapter is not a member of a team. If the adapter is a member of a team, remove the adapter from the team in Intel PROSet.

To uninstall the drivers and software from Windows Server, select **Intel(R) Network Connections** from **Programs and Features** in the Control Panel. To uninstall the adapter drivers, double-click on it or click the Remove button.

NOTE: The Device Manager should not be used to uninstall drivers. If the Device Manager is used to uninstall drivers, base drivers will not reinstall using the Modify option through Add/Remove Programs in the Control Panel.

Temporarily Disabling an Adapter

If you are testing your system to determine network faults, particularly in a multi-adapter environment, it is recommended that you temporarily disable the adapter.

1. From the Control Panel, double-click the **System** icon, click the **Hardware** tab, and click **Device Manager**.
2. Right-click the icon of adapter you want to disable, then click **Disable**.
3. Click **Yes** on the confirmation dialog box.

To enable the adapter, right-click its icon, then click **Enable**.

NOTE: You can also disable an adapter by right-clicking its icon in the **Network Connections** control panel and selecting **Disable**.

Replacing an Adapter

After installing an adapter in a specific slot, Windows treats any other adapter of the same type as a new adapter. Also, if you remove the installed adapter and insert it into a different slot, Windows recognizes it as a new adapter. Make sure that you follow the instructions below carefully.

1. Open Intel PROSet.
2. If the adapter is part of a team remove the adapter from the team.
3. Shut down the server and unplug the power cable.
4. Disconnect the network cable from the adapter.

5. Open the case and remove the adapter.
6. Insert the replacement adapter. (Use the same slot, otherwise Windows assumes that there is a new adapter.)
7. Reconnect the network cable.
8. Close the case, reattach the power cable, and power-up the server.
9. Open Intel PROSet and check to see that the adapter is available.
10. If the former adapter was part of a team, follow the instructions in Configuring ANS Teams to add the new adapter to the team.
11. If the former adapter was tagged with a VLAN, follow the instructions in Creating IEEE VLANs to tag the new adapter.

Removing an Adapter

Before physically removing an adapter from the system, be sure to complete these steps:

1. Use Intel PROSet to remove the adapter from any team or VLAN.
2. Uninstall the adapter drivers.

After you have completed these steps, power down the system, unplug the power cable and remove the adapter.

Using Advanced Features

You must use Intel PROSet to configure advanced features such as teaming or VLANs. Settings can be configured under the Intel PROSet for Windows Device Manager's Advanced tab. Some settings can also be configured using the Device Manager's adapter properties dialog box.

Using Intel® PROSet for Windows* Device Manager

Overview

Intel® PROSet for Windows* Device Manager is an extension to the Windows Device Manager. When you install the Intel PROSet software, additional tabs are automatically added to the supported Intel adapters in Device Manager. These features allow you to test and configure Intel wired network adapters.

Installing Intel PROSet for Windows Device Manager

Intel PROSet for Windows Device Manager is installed with the same process used to install drivers.

NOTES:

- You must have administrator rights to install or use Intel PROSet for Windows Device Manager.
- Upgrading PROSet for Windows Device Manager may take a few minutes.

1. On the autorun, click **Install Base Drivers and Software**.

   NOTE: You can also run setup64.exe from the files downloaded from Customer Support.

2. Proceed with the installation wizard until the **Custom Setup** page appears.
3. Select the features to install.
4. Follow the instructions to complete the installation.

If Intel PROSet for Windows Device Manager was installed without ANS support, you can install support by clicking **Install Base Drivers and Software** on the autorun, or running setup64.exe, and then selecting the **Modify** option when prompted. From the Intel® Network Connections window, select **Advanced Network Services** then click Next to continue with the installation wizard.

Using Intel PROSet for Windows Device Manager

The main Intel PROSet for Windows Device Manager window is similar to the illustration below. For more information about features on the custom Intel tabs, see the online help, which is integrated into the Properties dialog.

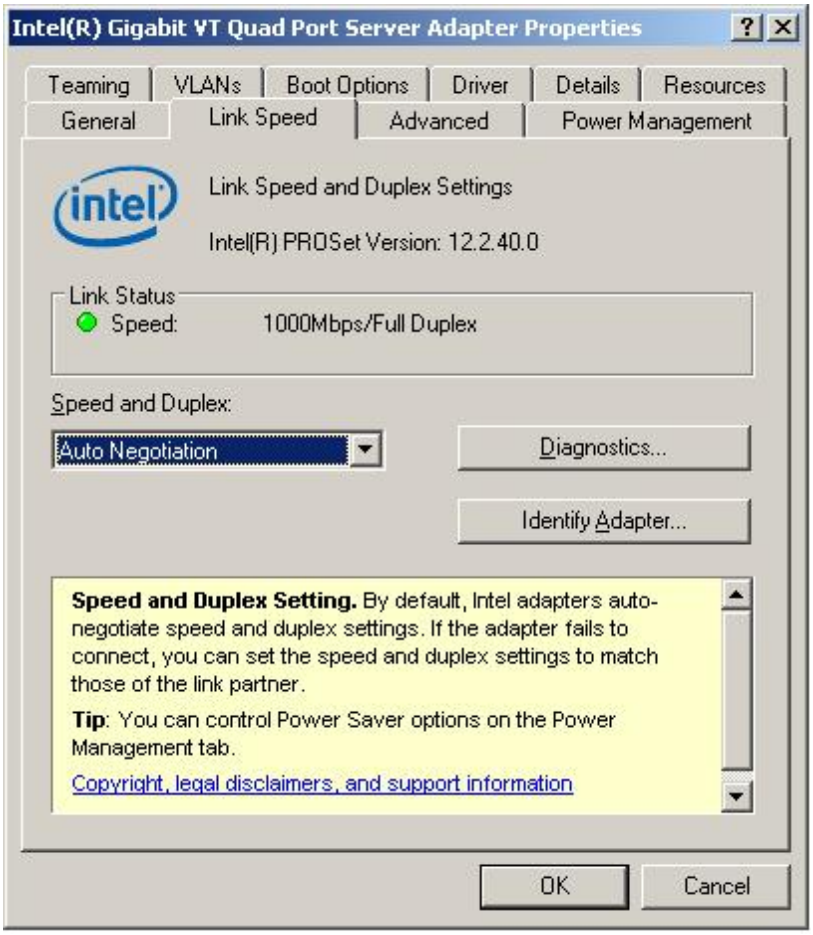

The **Link Speed** tab allows you to change the adapter's speed and duplex setting, run diagnostics, and use the identify adapter feature.

The **Advanced** tab allows you to change the adapter's advanced settings. These settings will vary on the type and model of adapter.

The **Teaming** tab allows you to create, modify, and delete adapter teams. You must install Advanced Network Services in order to see this tab and use the feature. See Installing Intel PROSet for Windows Device Manager for more information.

The **VLANs** tab allows you to create, modify, and delete VLANs. You must install Advanced Network Services in order to see this tab and use the feature. See Installing Intel PROSet for Windows Device Manager for more information.

The **Boot Options** tab allows you to configure Intel Boot Agent settings for the adapter.

NOTE: This tab will not appear if the Boot Agent has not been enabled on the adapter.

The **Power Management** tab allows you to configure power consumption settings for the adapter.

Configuring ANS Teams

Advanced Network Services (ANS) Teaming, a feature of the Advanced Network Services component, lets you take advantage of multiple adapters in a system by grouping them together. ANS teaming can use features like fault tolerance and load balancing to increase throughput and reliability.

Before you can set up ANS teaming in Windows*, you must install Intel® PROSet software. See Installing Intel PROSet for Windows Device Manager for more information.

NOTES:

- NLB will not work when Receive Load Balancing (RLB) is enabled. This occurs because NLB and iANS both attempt to set the server's multicast MAC address, resulting in an ARP table mismatch.
- Teaming with the Intel® 10 Gigabit AF DA Dual Port Server Adapter is only supported with similar

adapter types and models or with switches using a Direct Attach connection.

Creating a team

1. Launch Windows Device Manager
2. Expand **Network Adapters**.
3. Double-click on one of the adapters that will be a member of the team.
   The adapter properties dialog box appears.
4. Click the **Teaming** tab.
5. Click **Team with other adapters**.
6. Click **New Team**.
7. Type a name for the team, then click **Next**.
8. Click the checkbox of any adapter you want to include in the team, then click **Next**.
9. Select a teaming mode, then click **Next**.
   For more information on team types, see Set Up Adapter Teaming.
10. Click **Finish**.

The Team Properties window appears, showing team properties and settings.

Once a team has been created, it appears in the Network Adapters category in the Computer Management window as a virtual adapter. The team name also precedes the adapter name of any adapter that is a member of the team.

NOTE: If you want to set up VLANs on a team, you must first create the team.

Adding or Removing an Adapter from an Existing Team

NOTE: A team member should be removed from the team with link down. See the Configuration Notes in Adapter Teaming for more information.>

1. Open the Team Properties dialog box by double-clicking on a team listing in the Computer Management window.
2. Click the **Settings** tab.
3. Click **Modify Team**, then click the **Adapters** tab.
4. Select the adapters that will be members of the team.
   - Click the checkbox of any adapter that you want to add to the team.
   - Clear the checkbox of any adapter that you want to remove from the team.
5. Click **OK**.

Renaming a Team

1. Open the Team Properties dialog box by double-clicking on a team listing in the Computer Management window.
2. Click the **Settings** tab.
3. Click **Modify Team**, then click the **Name** tab.
4. Type a new team name, then click **OK**.

Removing a Team

1. Open the Team Properties dialog box by double-clicking on a team listing in the Computer Management window.
2. Click the **Settings** tab.
3. Select the team you want to remove, then click **Remove Team**.
4. Click **Yes** when prompted.

NOTE: If you defined a VLAN or QoS Prioritization on an adapter joining a team, you may have to redefine it when it is returned to a stand-alone mode.

Configuring IEEE VLANs

Before you can set up VLANs in Windows*, you must install Intel® PROSet software. See Installing Intel PROSet for Windows Device Manager for more information.

A maximum of 64 VLANs can be used on a server.

CAUTION:

- **VLANs cannot be used on teams that contain non-Intel network adapters**
- **Use Intel PROSet to add or remove a VLAN. Do not use the Network and Dial-up Connections dialog box to enable or disable VLANs. Otherwise, the VLAN driver may not be correctly enabled or disabled.**

NOTES:

- If you will be using both teaming and VLANs, be sure to set up teaming first.
- If you change a setting under the Advanced tab for one VLAN, it changes the settings for all VLANs using that port.

Setting Up an IEEE tagged VLAN

1. In the adapter properties window, click the **VLANs** tab.
2. Click **New**.
3. Type a name and ID number for the VLAN you are creating.
   The VLAN ID must match the VLAN ID on the switch. Valid ID range is from 1-4094, though your switch might not support this many IDs. The VLAN Name is for information only and does not have to match the name on the switch. The VLAN Name is limited to 256 characters.
   NOTE: VLAN IDs 0 and 1 are often reserved for other uses.
4. Click **OK**.

The VLAN entry will appear under Network Adapters in the Computer Management window.

Complete these steps for each adapter you want to add to a VLAN.

NOTE: If you configure a team to use VLANs, the team object icon in the Network Connections Panel will indicate that the team is disconnected. You will not be able to make any TCP/IP changes, such as changing an IP address or subnet mask. You will, however, be able to configure the team (add or remove team members, change team type, etc.) through Device Manager.

Setting Up an Untagged VLAN

You can set up only one untagged VLAN per adapter or team.

NOTE: An untagged VLAN cannot be created unless at least one tagged VLAN already exists.

1. In the adapter properties window, click the **VLANs** tab.
2. Click **New**.
3. Check the **Untagged VLAN** box.
4. Type a name for the VLAN you are creating.
   The VLAN name is for information only and does not have to match the name on the switch. It is limited to 256 characters.
5. Click **OK**.

Removing a VLAN

1. On the **VLANs** tab, select the VLAN you want to remove.
2. Click **Remove**.
3. Click **Yes** to confirm.

Removing Phantom Teams and Phantom VLANs

If you physically remove all adapters that are part of a team or VLAN from the system without removing them via the Device Manager first, a phantom team or phantom VLAN will appear in Device Manager. There are two methods to remove the phantom team or phantom VLAN.

Removing the Phantom Team or Phantom VLAN through the Device Manager

Follow these instructions to remove a phantom team or phantom VLAN from the Device Manager:

1. In the Device Manager, double-click on the phantom team or phantom VLAN.
2. Click the Settings tab.
3. Select Remove Team or Remove VLAN.

Removing the Phantom Team or Phantom VLAN using the savresdx.vbs Script

For Windows Server, the savresdx.vbs script is located in the driver update package in the WMI directory of the appropriate Windows folder. From the DOS command box type: "`cscript savresdx.vbs removephantoms`".

Preventing the Creation of Phantom Devices

To prevent the creation of phantom devices, make sure you perform these steps before physically removing an adapter from the system:

1. Remove the adapter from any teams using the Settings tab on the team properties dialog box.
2. Remove any VLANs from the adapter using the VLANs tab on the adapter properties dialog box.
3. Uninstall the adapter from Device Manager.

You do not need to follow these steps in hot-replace scenarios.

Removing Intel PROSet for Windows Device Manager

To uninstall the extensions to Windows Device Manager provided by Intel PROSet for Windows Device Manager, select **Intel(R) PRO Network Connections** from **Programs and Features** in the Control Panel.

NOTES:

- This process removes all Intel PRO adapter drivers and software.
- It is suggested that you uninstall VLANs and teams before removing adapters.
- The setup -u can also be used from the command line to remove Intel PROSet.

Configuring with IntelNetCmdlets Module for Windows PowerShell*

The IntelNetCmdlets module for Windows PowerShell contains several cmdlets that allow you to configure and manage the Intel® Ethernet Adapters and devices present in your system. For a complete list of these cmdlets and their descriptions, type **get-help IntelNetCmdlets** at the Windows PowerShell prompt. For detailed usage information for each cmdlet, type **get-help <*cmdlet_name*>** at the Windows PowerShell prompt.

NOTE: Online help (get-help -online) is not supported.

Install the IntelNetCmdlets module by checking the Windows PowerShell Module checkbox during the driver and PROSet installation process. Then use the Import-Module cmdlet to import the new cmdlets. You may need to restart Windows PowerShell to access the newly imported cmdlets.

To use the Import-Module cmdlet, you must specify the path and the name of the module. For example:

```
PS c:\> Import-Module -Name "C:\Program Files\Intel\IntelNetCmdlets"
```

See Microsoft TechNet for more information about the Import-Module cmdlet.

System requirements for using IntelNetCmdlets:

- Microsoft* Windows PowerShell* version 2.0
- .NET version 2.0

Configuring SR-IOV for improved network security

In a virtualized environment, on Intel® Server Adapters that support SR-IOV, the virtual function (VF) may be subject to malicious behavior. Software-generated frames are not expected and can throttle traffic between the host and the virtual switch, reducing performance. To resolve this issue, configure all SR-IOV enabled ports for VLAN tagging. This configuration allows unexpected, and potentially malicious, frames to be dropped.

Changing Intel PROSet Settings via Microsoft* Windows PowerShell*

You can use the IntelNetCmdlets module for Windows PowerShell to change most Intel PROSet settings.

NOTE: If an adapter is bound to an ANS team, do not change settings using the Set–NetAdapterAdvanceProperty cmdlet from Windows PowerShell*, or any other cmdlet not provided by Intel. Doing so may cause the team to stop using that adapter to pass traffic. You may see this as reduced performance or the adapter being disabled in the ANS team. You can resolve this issue by changing the setting back to its previous state, or by removing the adapter from the ANS team and then adding it back.

Advanced Settings for Windows* Drivers

The settings listed on Intel PROSet for Windows Device Manager's **Advanced** tab allow you to customize how the adapter handles QoS packet tagging, Jumbo Packets, Offloading, and other capabilities. Some of the following features might not be available depending on the operating system you are running, the specific adapters installed, and the specific platform you are using.

Gigabit Master Slave Mode

Determines whether the adapter or link partner is designated as the master. The other device is designated as the slave. By default, the IEEE 802.3ab specification defines how conflicts are handled. Multi-port devices such as switches have higher priority over single port devices and are assigned as the master. If both devices are multi-port devices, the one with higher seed bits becomes the master. This default setting is called "Hardware Default."

NOTE: In most scenarios, it is recommended to keep the default value of this feature.

Setting this to either "Force Master Mode" or "Force Slave Mode" overrides the hardware default.

Default	Auto Detect
Range	• Force Master Mode • Force Slave Mode • Auto Detect

NOTE: Some multi-port devices may be forced to Master Mode. If the adapter is connected to such a device and is configured to "Force Master Mode," link is not established.

Jumbo Frames

Enables or disables Jumbo Packet capability. The standard Ethernet frame size about 1514 bytes, while Jumbo Packets are larger than this. Jumbo Packets can increase throughput and decrease CPU utilization. However, additional latency may be introduced.

Enable Jumbo Packets only if ALL devices across the network support them and are configured to use the same frame size. When setting up Jumbo Packets on other network devices, be aware that network devices calculate Jumbo Packet sizes differently. Some devices include the frame size in the header information while others do not. Intel adapters do not include frame size in the header information.

Jumbo Packets can be implemented simultaneously with VLANs and teaming. If a team contains one or more non-Intel adapters, the Jumbo Packets feature for the team is not supported. Before adding a non-Intel adapter to a team, make sure that you disable Jumbo Packets for all non-Intel adapters using the software shipped with the adapter.

Restrictions

- Jumbo frames are not supported in multi-vendor team configurations.
- Supported protocols are limited to IP (TCP, UDP).
- Jumbo frames require compatible switch connections that forward Jumbo Frames. Contact your switch vendor for more information.
- When standard-sized Ethernet frames (64 to 1518 bytes) are used, there is no benefit to configuring Jumbo Frames.
- The Jumbo Packets setting on the switch must be set to at least 8 bytes larger than the adapter setting for Microsoft Windows operating systems, and at least 22 bytes larger for all other operating systems.

Default	Disabled
Range	Disabled (1514), 4088, or 9014 bytes. (Set the switch 4 bytes higher for CRC, plus 4 bytes if using VLANs.)

NOTES:

- Jumbo Packets are supported at 10 Gbps and 1 Gbps only. Using Jumbo Packets at 10 or 100 Mbps may result in poor performance or loss of link.
- End-to-end hardware must support this capability; otherwise, packets will be dropped.
- Intel adapters that support Jumbo Packets have a frame size limit of 9238 bytes, with a corresponding MTU size limit of 9216 bytes.

Locally Administered Address

Overrides the initial MAC address with a user-assigned MAC address. To enter a new network address, type a 12-digit hexadecimal number in this box.

Default	None
Range	0000 0000 0001 - FFFF FFFF FFFD Exceptions: • Do not use a multicast address (Least Significant Bit of the high byte = 1). For example, in the address 0Y123456789A, "Y" cannot be an odd number. (Y must be 0, 2, 4, 6, 8, A, C, or E.) • Do not use all zeros or all Fs. If you do not enter an address, the address is the original network address of the adapter. For example, Multicast: 0123 4567 8999 Broadcast: FFFF FFFF FFFF Unicast (legal): 0070 4567 8999

NOTE: In a team, Intel PROSet uses either:

- The primary adapter's permanent MAC address if the team does not have an LAA configured, or
- The team's LAA if the team has an LAA configured.

Intel PROSet does not use an adapter's LAA if the adapter is the primary adapter in a team and the team has an LAA.

Log Link State Event

This setting is used to enable/disable the logging of link state changes. If enabled, a link up change event or a link down change event generates a message that is displayed in the system event logger. This message contains the link's speed and duplex. Administrators view the event message from the system event log.

The following events are logged.

- The link is up.
- The link is down.
- Mismatch in duplex.
- Spanning Tree Protocol detected.

Default	Enabled
Range	Enabled, Disabled

Priority & VLAN Tagging

Enables the adapter to offload the insertion and removal of priority and VLAN tags for transmit and receive.

Default	Priority & VLAN Enabled
Range	• Priority & VLAN Disabled • Priority Enabled • VLAN Enabled • Priority & VLAN Enabled

Receive Side Scaling

When Receive Side Scaling (RSS) is enabled, all of the receive data processing for a particular TCP connection is shared across multiple processors or processor cores. Without RSS all of the processing is performed by a single processor, resulting in less efficient system cache utilization. RSS can be enabled for a LAN or for FCoE. In the first case, it is called "LAN RSS". In the second, it is called "FCoE RSS".

LAN RSS

LAN RSS applies to a particular TCP connection.

NOTE: This setting has no effect if your system has only one processing unit.

LAN RSS Configuration

RSS is enabled on the **Advanced** tab of the adapter property sheet. If your adapter does not support RSS, or if the SNP or SP2 is not installed, the RSS setting will not be displayed. If RSS is supported in your system environment, the following will be displayed:

- **Port NUMA Node**. This is the NUMA node number of a device.
- **Starting RSS CPU**. This setting allows you to set the p referred starting RSS processor. Change this setting if the current processor is dedicated to other processes. The setting range is from 0 to the number of logical CPUs - 1. In Server 2008 R2, RSS will only use CPUs in group 0 (CPUs 0 through 63).
- **Max number of RSS CPU**. This setting allows you to set the maximum number of CPUs assigned to an adapter and is primarily used in a Hyper-V environment. By decreasing this setting in a Hyper-V environment, the total number of interrupts is reduced which lowers CPU utilization. The default is 8 for Gigabit adapters and 16 for 10 Gigabit adapters.
- **Preferred NUMA Node**. This setting allows you to choose the preferred NUMA (Non-Uniform Memory Access) node to be used for memory allocations made by the network adapter. In addition the system will attempt to use the CPUs from the preferred NUMA node first for the purposes of RSS. On NUMA platforms, memory access latency is dependent on the memory location. Allocation of memory from the closest node helps improve performance. The Windows Task Manager shows the NUMA Node ID for each processor.

  NOTE: This setting only affects NUMA systems. It will have no effect on non-NUMA systems.

- **Receive Side Scaling Queues**. This setting configures the number of RSS queues, which determine the space to buffer transactions between the network adapter and CPU(s).

Default	2 queues for the Intel® 10 Gigabit Server Adapters
Range	• 1 queue is used when low CPU utilization is required.

	• 2 queues are used when good throughput and low CPU utilization are required. • 4 queues are used for applications that demand maximum throughput and transactions per second. • 8 and 16 queues are supported on the Intel® 82598-based and 82599-based adapters. **NOTES:** • The 8 and 16 queues are only available when PROSet for Windows Device Manager is installed. If PROSet is not installed, only 4 queues are available. • Using 8 or more queues will require the system to reboot. **NOTE:** Not all settings are available on all adapters.

LAN RSS and Teaming

- If RSS is not enabled for all adapters in a team, RSS will be disabled for the team.
- If an adapter that does not support RSS is added to a team, RSS will be disabled for the team.
- If you create a multi-vendor team, you must manually verify that the RSS settings for all adapters in the team are the same.

FCoE RSS

If FCoE is installed, FCoE RSS is enabled and applies to FCoE receive processing that is shared across processor cores.

FCoE RSS Configuration

If your adapter supports FCoE RSS, the following configuration settings can be viewed and changed on the base driver Advanced Performance tab:

- **FCoE NUMA Node Count**. This setting specifies the number of consecutive NUMA Nodes where the allocated FCoE queues will be evenly distributed.
- **FCoE Starting NUMA Node**. This setting specifies the NUMA node representing the first node within the FCoE NUMA Node Count.
- **FCoE Starting Core Offset**. This setting specifies the offset to the first NUMA Node CPU core that will be assigned to FCoE queue.
- **FCoE Port NUMA Node**. This setting in an indication from the platform of the optimal closest NUMA Node to the physical port, if available. This setting is read-only and cannot be configured.

Performance Tuning

The Intel Network Controller provides a new set of advanced FCoE performance tuning options. These options will direct how FCoE transmit/receive queues are allocated in NUMA platforms. Specifically, they direct what target set of NUMA node CPUs can be selected from to assign individual queue affinity. Selecting a specific CPU has two main effects:

- It sets the desired interrupt location for processing queue packet indications.
- It sets the relative locality of the queue to available memory.

As indicated, these are intended as advanced tuning options for those platform managers attempting to maximize system performance. They are generally expected to be used to maximize performance for multi-port platform configurations. Since all ports share the same default installation directives (the .inf file, etc.), the FCoE queues for every port will be associated with the same set of NUMA CPUs which may result in CPU contention.

The software exporting these tuning options defines a NUMA Node to be equivalent to an individual processor (socket). Platform ACPI information presented by the BIOS to the operating system helps define the relation of PCI devices to individual processors. However, this detail is not currently reliably provided in all platforms. Therefore, using the tuning options may produce unexpected results. Consistent or predictable results when using the performance options cannot be guaranteed.

The performance tuning options are listed in the LAN RSS Configuration section.

Example 1: A platform with two physical sockets, each socket processor providing 8 core CPUs (16 when hyper threading is enabled), and a dual port Intel adapter with FCoE enabled.

By default 8 FCoE queues will be allocated per NIC port. Also, by default the first (non-hyper thread) CPU cores of the first processor will be assigned affinity to these queues resulting in the allocation model pictured below. In this scenario, both ports would be competing for CPU cycles from the same set of CPUs on socket 0.

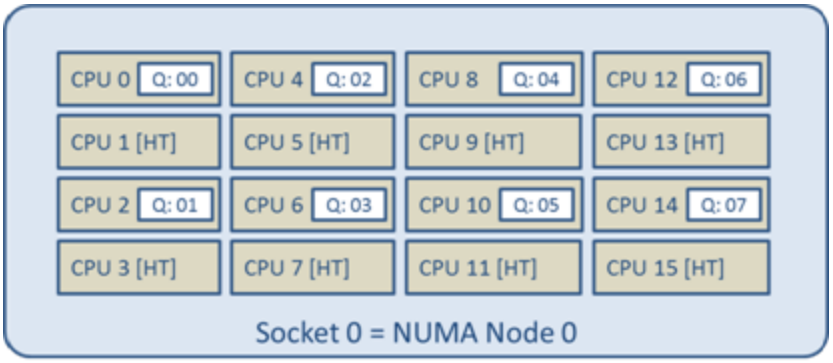

Socket Queue to CPU Allocation

Using performance tuning options, the association of the FCoE queues for the second port can be directed to a different non-competing set of CPU cores. The following settings would direct SW to use CPUs on the other processor socket:

- FCoE NUMA Node Count = 1: Assign queues to cores from a single NUMA node (or processor socket).
- FCoE Starting NUMA Node = 1: Use CPU cores from the second NUMA node (or processor socket) in the system.
- FCoE Starting Core Offset = 0: SW will start at the first CPU core of the NUMA node (or processor socket).

The following settings would direct SW to use a different set of CPUs on the same processor socket. This assumes a processor that supports 16 non-hyperthreading cores.

- FCoE NUMA Node Count = 1
- FCoE Starting NUMA Node = 0
- FCoE Starting Core Offset = 8

Example 2: Using one or more ports with queues allocated across multiple NUMA nodes. In this case, for each NIC port the FCoE NUMA Node Count is set to that number of NUMA nodes. By default the queues will be allocated evenly from each NUMA node:

- FCoE NUMA Node Count = 2
- FCoE Starting NUMA Node = 0
- FCoE Starting Core Offset = 0

Example 3: The display shows FCoE Port NUMA Node setting is 2 for a given adapter port. This is a read-only indication from SW that the optimal nearest NUMA node to the PCI device is the third logical NUMA node in the system. By default SW has allocated that port's queues to NUMA node 0. The following settings would direct SW to use CPUs on the optimal processor socket:

- FCoE NUMA Node Count = 1
- FCoE Starting NUMA Node = 2
- FCoE Starting Core Offset = 0

This example highlights the fact that platform architectures can vary in the number of PCI buses and where they are attached. The figures below show two simplified platform architectures. The first is the older common FSB style architecture in which multiple CPUs share access to a single MCH and/or ESB that provides PCI bus and memory connectivity. The second is a more recent architecture in which multiple CPU processors are interconnected via QPI, and each processor itself supports integrated MCH and PCI connectivity directly.

There is a perceived advantage in keeping the allocation of port objects, such as queues, as close as possible to the NUMA node or collection of CPUs where it would most likely be accessed. If the port queues are using CPUs and memory from one socket when the PCI device is actually hanging off of another socket, the result may be undesirable QPI processor-to-processor bus bandwidth being consumed. It is important to understand the platform architecture when using these performance options.

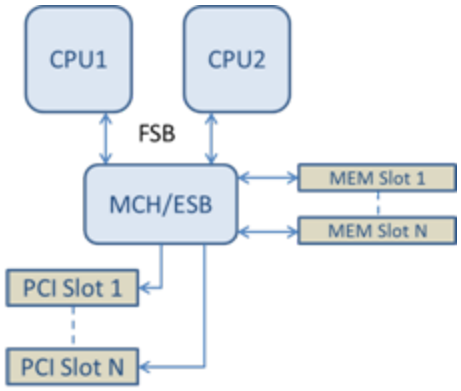

Shared Single Root PCI/Memory Architecture

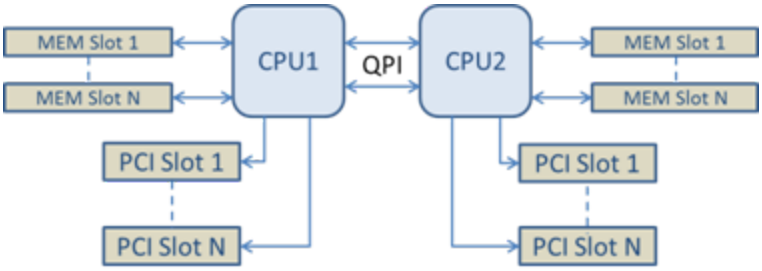

Distributed Multi-Root PCI/Memory Architecture

Example 4: The number of available NUMA node CPUs is not sufficient for queue allocation. If your platform has a processor that does not support an even power of 2 CPUs (for example, it supports 6 cores), then during queue allocation if SW runs out of CPUs on one socket it will by default reduce the number of queues to a power of 2 until allocation is achieved. For example, if there is a 6 core processor being used, the SW will only allocate 4 FCoE queues if there only a single NUMA node. If there are multiple NUMA nodes, the NUMA node count can be changed to a value greater than or equal to 2 in order to have all 8 queues created.

Determining Active Queue Location

The user of these performance options will want to determine the affinity of FCoE queues to CPUs in order to verify their actual effect on queue allocation. This is easily done by using a small packet workload and an I/O application such as IoMeter. IoMeter monitors the CPU utilization of each CPU using the built-in performance monitor provided by the operating system. The CPUs supporting the queue activity should stand out. They should be the first non-hyper thread CPUs available on the processor unless the allocation is specifically directed to be shifted via the performance options discussed above.

To make the locality of the FCoE queues even more obvious, the application affinity can be assigned to an isolated set of CPUs on the same or another processor socket. For example, the IoMeter application can be set to run only on a finite number of hyper thread CPUs on any processor. If the performance options have been set to direct queue allocation on a specific NUMA node, the application affinity can be set to a different NUMA node. The FCoE queues should not move and the activity should remain on those CPUs even though the application CPU activity moves to the other processor CPUs selected.

Wait for Link

Determines whether the driver waits for auto-negotiation to be successful before reporting the link state. If this feature is off, the driver does not wait for auto-negotiation. If the feature is on, the driver does wait for auto-negotiation.

If this feature is on, and the speed is not set to auto-negotiation, the driver will wait for a short time for link to complete before reporting the link state.

If the feature is set to **Auto Detect**, this feature is automatically set to **On** or **Off** depending on speed and adapter type when the driver is installed. The setting is:

- Off for copper Intel gigabit adapters with a speed of "Auto".
- On for copper Intel gigabit adapters with a forced speed and duplex.
- On for fiber Intel gigabit adapters with a speed of "Auto".

Default	Auto Detect
Range	• On • Off • Auto Detect

Thermal Monitoring

Adapters and network controllers based on the Intel® Ethernet Controller I350 (and later controllers) can display temperature data and automatically reduce the link speed if the controller temperature gets too hot.

NOTE: This feature is enabled and configured by the equipment manufacturer. It is not available on all adapters and network controllers. There are no user configurable settings.

Monitoring and Reporting

Temperature information is displayed on the **Link** tab in Intel® PROSet for Windows* Device Manger. There are three possible conditions:

- Temperature: Normal
  Indicates normal operation.
- Temperature: Overheated, Link Reduced
  Indicates that the device has reduced link speed to lower power consumption and heat.
- Temperature: Overheated, Adapter Stopped
  Indicates that the device is too hot and has stopped passing traffic so it is not damaged.

If either of the overheated events occur, the device driver writes a message to the system event log.

Performance Options

Adaptive Inter-Frame Spacing

Compensates for excessive Ethernet packet collisions on the network.

The default setting works best for most computers and networks. By enabling this feature, the network adapter dynamically adapts to the network traffic conditions. However, in some rare cases you might obtain better performance by disabling this feature. This setting forces a static gap between packets.

Default	Disabled
Range	• Enabled • Disabled

Direct Memory Access (DMA) Coalescing

DMA (Direct Memory Access) allows the network device to move packet data directly to the system's memory, reducing CPU utilization. However, the frequency and random intervals at which packets arrive do not allow the system to enter a lower power state. DMA Coalescing allows the NIC to collect packets before it initiates a DMA event. This may increase network latency but also increases the chances that the system will consume less energy. Adapters and network devices based on the Intel® Ethernet Controller I350 (and later controllers) support DMA Coalescing.

Higher DMA Coalescing values result in more energy saved but may increase your system's network latency. If you enable DMA Coalescing, you should also set the Interrupt Moderation Rate to 'Minimal'. This minimizes the latency impact imposed by DMA Coalescing and results in better peak network throughput performance. You must enable DMA Coalescing on all active ports in the system. You may not gain any energy savings if it is enabled only on some of the ports in your system. There are also several BIOS, platform, and application settings that will affect your potential energy savings. A white paper containing information on how to best configure your platform is available on the Intel website.

Flow Control

Enables adapters to more effectively regulate traffic. Adapters generate flow control frames when their receive queues reach a pre-defined limit. Generating flow control frames signals the transmitter to slow transmission. Adapters respond to flow control frames by pausing packet transmission for the time specified in the flow control frame.

By enabling adapters to adjust packet transmission, flow control helps prevent dropped packets.

NOTES:

- For adapters to benefit from this feature, link partners must support flow control frames.
- When an adapter is running in NPar mode, Flow Control is limited to the root partition of each port.

Default	RX & TX Enabled
Range	• Disabled • RX Enabled

	• TX Enabled • RX & TX Enabled

Interrupt Moderation Rate

Sets the Interrupt Throttle Rate (ITR). This setting moderates the rate at which Transmit and Receive interrupts are generated.

When an event such as packet receiving occurs, the adapter generates an interrupt. The interrupt interrupts the CPU and any application running at the time, and calls on the driver to handle the packet. At greater link speeds, more interrupts are created, and CPU rates also increase. This results in poor system performance. When you use a higher ITR setting, the interrupt rate is lower and the result is better CPU performance.

NOTE: A higher ITR rate also means that the driver has more latency in handling packets. If the adapter is handling many small packets, it is better to lower the ITR so that the driver can be more responsive to incoming and outgoing packets.

Altering this setting may improve traffic throughput for certain network and system configurations, however the default setting is optimal for common network and system configurations. Do not change this setting without verifying that the desired change will have a positive effect on network performance.

Default	Adaptive
Range	• Adaptive • Extreme • High • Medium • Low • Minimal • Off

Low Latency Interrupts

LLI enables the network device to bypass the configured interrupt moderation scheme based on the type of data being received. It configures which arriving TCP packets trigger an immediate interrupt, enabling the system to handle the packet more quickly. Reduced data latency enables some applications to gain faster access to network data.

NOTE: When LLI is enabled, system CPU utilization may increase.

LLI can be used for data packets containing a TCP PSH flag in the header or for specified TCP ports.

- **Packets with TCP PSH Flag** - Any incoming packet with the TCP PSH flag will trigger an immediate interrupt. The PSH flag is set by the sending device.
- **TCP Ports** - Every packet received on the specified ports will trigger an immediate interrupt. Up to eight ports may be specified.

Default	Disabled
Range	• Disabled • PSH Flag-Based • Port-Based

Receive Buffers

Defines the number of Receive Buffers, which are data segments. They are allocated in the host memory and used to store the received packets. Each received packet requires at least one Receive Buffer, and each buffer uses 2KB of memory.

You might choose to increase the number of Receive Buffers if you notice a significant decrease in the performance of received traffic. If receive performance is not an issue, use the default setting appropriate to the adapter.

Default	512, for the 10 Gigabit Server Adapters. 256, for all other adapters depending on the features selected.
Range	128-4096, in intervals of 64, for the 10 Gigabit Server Adapters. 80-2048, in intervals of 8, for all other adapters.
Recommended Value	Teamed adapter: 256 Using IPSec and/or multiple features: 352

Transmit Buffers

Defines the number of Transmit Buffers, which are data segments that enable the adapter to track transmit packets in the system memory. Depending on the size of the packet, each transmit packet requires one or more Transmit Buffers.

You might choose to increase the number of Transmit Buffers if you notice a possible problem with transmit performance. Although increasing the number of Transmit Buffers can enhance transmit performance, Transmit Buffers do consume system memory. If transmit performance is not an issue, use the default setting. This default setting varies with the type of adapter.

View the Adapter Specifications topic for help identifying your adapter.

Default	512, depending on the requirements of the adapter
Range	128-16384, in intervals of 64, for 10 Gigabit Server Adapters. 80-2048, in intervals of 8, for all other adapters.

Performance Profile

Performance Profiles are supported on Intel® 10GbE adapters and allow you to quickly optimize the performance of your Intel® Ethernet Adapter. Selecting a performance profile will automatically adjust some Advanced Settings to their optimum setting for the selected application. For example, a standard server has optimal performance with only two RSS (Receive-Side Scaling) queues, but a web server requires more RSS queues for better scalability.

You must install Intel® PROSet for Windows Device Manager to use Performance profiles. Profiles are selected on the Advanced tab of the adapter's property sheet.

Profiles	• Standard Server – This profile is optimized for typical servers. • Web Server – This profile is optimized for IIS and HTTP-based web servers. • Virtualization Server – This profile is optimized for Microsoft's Hyper-V virtualization environment. • Storage Server – This profile is optimized for Fibre Channel over Ethernet or for iSCSI over DCB performance. Selecting this profile will disable SR-IOV and VMQ. • Storage + Virtualization – This profile is optimized for a combination of storage and virtualization requirements. • Low Latency – This profile is optimized to minimize network latency.

NOTES:

- Not all options are available on all adapter/operating system combinations.
- If you have selected the Virtualization Server profile or the Storage + Virtualization profile, and you uninstall the Hyper-V role, you should select a new profile.

Teaming Considerations

When you create a team with all members of the team supporting Performance Profiles, you will be asked which profile to use at the time of team creation. The profile will be synchronized across the team. If there is not a profile that is supported by all team members then the only option will be Use Current Settings. The team will be created normally. Adding an adapter to an existing team works in much the same way.

If you attempt to team an adapter that supports performance profiles with an adapter that doesn't, the profile on the supporting adapter will be set to Custom Settings and the team will be created normally.

TCP/IP Offloading Options

IPv4 Checksum Offload

This allows the adapter to compute the IPv4 checksum of incoming and outgoing packets. This feature enhances IPv4 receive and transmit performance and reduces CPU utilization.

With Offloading off, the operating system verifies the IPv4 checksum.

With Offloading on, the adapter completes the verification for the operating system.

Default	RX & TX Enabled
Range	• Disabled • RX Enabled • TX Enabled • RX & TX Enabled

Large Send Offload (IPv4 and IPv6)

Sets the adapter to offload the task of segmenting TCP messages into valid Ethernet frames. The maximum frame size limit for large send offload is set to 64,000 bytes.

Since the adapter hardware is able to complete data segmentation much faster than operating system software, this feature may improve transmission performance. In addition, the adapter uses fewer CPU resources.

Default	Enabled
Range	• Enabled • Disabled

TCP Checksum Offload (IPv4 and IPv6)

Allows the adapter to verify the TCP checksum of incoming packets and compute the TCP checksum of outgoing packets. This feature enhances receive and transmit performance and reduces CPU utilization.

With Offloading off, the operating system verifies the TCP checksum.

With Offloading on, the adapter completes the verification for the operating system.

Default	RX & TX Enabled
Range	• Disabled • RX Enabled • TX Enabled • RX & TX Enabled

UDP Checksum Offload (IPv4 and IPv6)

Allows the adapter to verify the UDP checksum of incoming packets and compute the UDP checksum of outgoing packets. This feature enhances receive and transmit performance and reduces CPU utilization.

With Offloading off, the operating system verifies the UDP checksum.

With Offloading on, the adapter completes the verification for the operating system.

Default	RX & TX Enabled
Range	• Disabled

	• RX Enabled • TX Enabled • RX & TX Enabled

Power Management Settings for Windows* Drivers

The Intel® PROSet **Power Management** tab replaces the standard Microsoft* Windows* Power Management tab in Device Manager. It includes the Power Saver options that were previously included on the Advanced tab. The standard Windows power management functionality is incorporated on the Intel PROSet tab.

NOTES:

- The Intel® 10 Gigabit Network Adapters do not support power management.
- If your system has a Manageability Engine, the Link LED may stay lit even if WoL is disabled.
- When an adapter is running in NPar mode, Power Management is limited to the root partition of each port.

Power Saver Options

The Intel PROSet Power Management tab includes several settings that control the adapter's power consumption. For example, you can set the adapter to reduce its power consumption if the cable is disconnected.

Reduce Power if Cable Disconnected & Reduce Link Speed During Standby

Enables the adapter to reduce power consumption when the LAN cable is disconnected from the adapter and there is no link. When the adapter regains a valid link, adapter power usage returns to its normal state (full power usage).

The Hardware Default option is available on some adapters. If this option is selected, the feature is disabled or enabled based on the system hardware.

Default	The default varies with the operating system and adapter.
Range	The range varies with the operating system and adapter.

Energy Efficient Ethernet

The Energy Efficient Ethernet (EEE) feature allows a capable device to enter Low-Power Idle between bursts of network traffic. Both ends of a link must have EEE enabled for any power to be saved. Both ends of the link will resume full power when data needs to be transmitted. This transition may introduce a small amount of network latency.

NOTES:

- Both ends of the EEE link must automatically negotiate link speed.
- EEE is not supported at 10 Mbps.

Wake on LAN Options

The ability to remotely wake computers is an important development in computer management. This feature has evolved over the last few years from a simple remote power-on capability to a complex system interacting with a variety of device and operating system power states. More details are available here.

Microsoft Windows Server is ACPI-capable. Windows does not support waking from a power-off (S5) state, only from standby (S3) or hibernate (S4). When shutting down the system, these states shut down ACPI devices, including Intel adapters. This disarms the adapter's remote wake-up capability. However, in some ACPI-capable computers, the BIOS may have a setting that allows you to override the operating system and wake from an S5 state anyway. If there is no support for wake from S5 state in your BIOS settings, you are limited to Wake From Standby when using these operating systems in ACPI computers.

The Intel PROSet Power Management tab includes **Wake on Magic Packet** and **Wake on directed packet settings**. These control the type of packets that wake up the system from standby.

For some adapters, the Power Management tab in Intel PROSet includes a setting called **Wake on Magic Packet from power off state**. Enable this setting to explicitly allow wake-up with a Magic Packet* from shutdown under APM power management mode.

NOTE: To use the Wake on Directed Packet feature, WoL must first be enabled in the EEPROM using BootUtil.

WoL Supported Devices

All devices support Wake on LAN on all ports, with the following exceptions:

Gigabit Adapters	**Adapter Port(s) supporting WoL**
Intel® Gigabit 2P I350-t Adapter	port A
Intel® Gigabit 4P I350-t Adapter	port A
10 Gigabit Adapters	
Intel® Ethernet 10G 4P X540/I350 rNDC	both 10G ports
Intel® Ethernet 10G 4P X520/I350 rNDC	both 10G ports
Intel® Converged Network Adapter X710	port 1

Wake on Link Settings

Wakes the computer if the network connection establishes link while the computer is in standby mode. You can enable the feature, disable it, or let the operating system use its default.

NOTES:

- If a copper-based Intel adapter is advertising a speed of one gigabit only, this feature does not work because the adapter cannot identify a gigabit link at a D3 state.
- The network cable must be disconnected when entering into S3/S4 in order to wake the system up by link up event.

Default	Disabled
Range	Disabled OS Controlled Forced

Microsoft* Hyper-V* Overview

Microsoft* Hyper-V* makes it possible for one or more operating systems to run simultaneously on the same physical system as virtual machines. This allows you to consolidate several servers onto one system, even if they are running different operating systems. Intel® Network Adapters work with, and within, Microsoft Hyper-V virtual machines with their standard drivers and software.

NOTES:

- Some virtualization options are not available on some adapter/operating system combinations.
- The jumbo frame setting inside a virtual machine must be the same, or lower than, the setting on the physical port.
- See http://www.intel.com/technology/advanced_comm/virtualization.htm for more information on using Intel Network Adapters in virtualized environments.

Using Intel® Network Adapters in a Hyper-V Environment

When a Hyper-V Virtual NIC (VNIC) interface is created in the parent partition, the VNIC takes on the MAC address of the underlying physical NIC. The same is true when a VNIC is created on a team or VLAN. Since the VNIC uses the MAC address of the underlying interface, any operation that changes the MAC address of the interface (for example, setting LAA on the interface, changing the primary adapter on a team, etc.), will cause the VNIC to lose connectivity. In order to prevent this loss of connectivity, Intel® PROSet will not allow you to change settings that change the MAC address.

NOTES:

- If Fibre Channel over Ethernet (FCoE)/Data Center Bridging (DCB) is present on the port, configuring the device in Virtual Machine Queue (VMQ) + DCB mode reduces the number of VMQ VPorts available for guest OSes. This does not apply to Intel® Ethernet Controller X710 based devices.
- When sent from inside a virtual machine, LLDP and LACP packets may be a security risk. The Intel® Virtual Function driver blocks the transmission of such packets.
- The Virtualization setting on the Advanced tab of the adapter's Device Manager property sheet is not available if the Hyper-V role is not installed.
- ANS teaming of VF devices inside a Windows 2008 R2 guest running on an open source hypervisor is supported.

The Virtual Machine Switch

The virtual machine switch is part of the network I/O data path. It sits between the physical NIC and the virtual machine NICs and routes packets to the correct MAC address. Enabling Virtual Machine Queue (VMQ) offloading in Intel(R) PROSet will automatically enable VMQ in the virtual machine switch. For driver-only installations, you must manually enable VMQ in the virtual machine switch.

Using ANS VLANs

If you create ANS VLANs in the parent partition, and you then create a Hyper-V Virtual NIC interface on an ANS VLAN, then the Virtual NIC interface ***must*** have the same VLAN ID as the ANS VLAN. Using a different VLAN ID or not setting a VLAN ID on the Virtual NIC interface will result in loss of communication on that interface.

Virtual Switches bound to an ANS VLAN will have the same MAC address as the VLAN, which will have the same address as the underlying NIC or team. If you have several VLANs bound to a team and bind a virtual switch to each VLAN, all of the virtual switches will have the same MAC address. Clustering the virtual switches together will cause a network error in Microsoft's cluster validation tool. In some cases, ignoring this error will not impact the performance of the cluster. However, such a cluster is not supported by Microsoft. Using Device Manager to give each of the virtual switches a unique address will resolve the issue. See the Microsoft TechNet article *Configure MAC Address Spoofing for Virtual Network Adapters* for more information.

Virtual Machine Queues (VMQ) and SR-IOV cannot be enabled on a Hyper-V Virtual NIC interface bound to a VLAN configured using the VLANs tab in Windows Device Manager.

Using an ANS Team or VLAN as a Virtual NIC

If you want to use a team or VLAN as a virtual NIC you must follow these steps:

NOTES:

- This applies only to virtual NICs created on a team or VLAN. Virtual NICs created on a physical adapter do not require these steps.
- Receive Load Balancing (RLB) is not supported in Hyper-V. Disable RLB when using Hyper-V.

1. Use Intel® PROSet to create the team or VLAN.
2. Open the Network Control Panel.
3. Open the team or VLAN.
4. On the General Tab, uncheck all of the protocol bindings and click OK.
5. Create the virtual NIC. (If you check the "Allow management operating system to share the network adapter." box you can do the following step in the parent partition.)
6. Open the Network Control Panel for the Virtual NIC.
7. On the General Tab, check the protocol bindings that you desire.

Command Line for Microsoft Windows Server* Core

Microsoft Windows Server* Core does not have a GUI interface. If you want to use an ANS Team or VLAN as a Virtual NIC, you must use Microsoft* Windows PowerShell* to set up the configuration. Use Windows PowerShell to create the team or VLAN.

Virtual Machine Queue Offloading

Enabling VMQ offloading increases receive and transmit performance, as the adapter hardware is able to perform these tasks faster than the operating system. Offloading also frees up CPU resources. Filtering is based on MAC and/or VLAN filters. For devices that support it, VMQ offloading is enabled in the host partition on the adapter's Device Manager property sheet, under Virtualization on the Advanced Tab.

Each Intel® Ethernet Adapter has a pool of virtual ports that are split between the various features, such as VMQ Offloading, SR-IOV, Data Center Bridging (DCB), and Fibre Channel over Ethernet (FCoE). Increasing the number of virtual ports used for one feature decreases the number available for other features. On devices that support it, enabling DCB reduces the total pool available for other features to 32. Enabling FCoE further reduces the total pool to 24.

NOTE: This does not apply to devices based on the Intel® Ethernet X710 or XL710 controllers.

Intel PROSet displays the number of virtual ports available for virtual functions under Virtualization properties on the device's Advanced Tab. It also allows you to set how the available virtual ports are distributed between VMQ and SR-IOV.

Teaming Considerations

- If VMQ is not enabled for all adapters in a team, VMQ will be disabled for the team.
- If an adapter that does not support VMQ is added to a team, VMQ will be disabled for the team.
- Virtual NICs cannot be created on a team with Receive Load Balancing enabled. Receive Load Balancing is automatically disabled if you create a virtual NIC on a team.
- If a team is bound to a Hyper-V virtual NIC, you cannot change the Primary or Secondary adapter.

Virtual Machine Multiple Queues

Virtual Machine Multiple Queues (VMMQ) enables Receive Side Scaling (RSS) for virtual ports attached to a physical port. This allows RSS to be used with SR-IOV and inside a VMQ virtual machine, and offloads the RSS processing to the network adapter. RSS balances receive traffic across multiple CPUs or CPU cores. This setting has no effect if your system has only one processing unit.

SR-IOV (Single Root I/O Virtualization)

SR-IOV lets a single network port appear to be several virtual functions in a virtualized environment. If you have an SR-IOV capable NIC, each port on that NIC can assign a virtual function to several guest partitions. The virtual functions bypass the Virtual Machine Manager (VMM), allowing packet data to move directly to a guest partition's memory, resulting in higher throughput and lower CPU utilization. SR-IOV also allows you to move packet data directly to a guest partition's memory. SR-IOV support was added in Microsoft Windows Server 2012. See your operating system documentation for system requirements.

For devices that support it, SR-IOV is enabled in the host partition on the adapter's Device Manager property sheet, under Virtualization on the Advanced Tab. Some devices may need to have SR-IOV enabled in a preboot environment.

NOTES:

- **Configuring SR-IOV for improved network security:** In a virtualized environment, on Intel® Server Adapters that support SR-IOV, the virtual function (VF) may be subject to malicious behavior. Software-generated frames are not expected and can throttle traffic between the host and the virtual switch, reducing performance. To resolve this issue, configure all SR-IOV enabled ports for VLAN tagging. This configuration allows unexpected, and potentially malicious, frames to be dropped.
- You must enable VMQ for SR-IOV to function.
- SR-IOV is not supported with ANS teams.
- VMWare ESXi does not support SR-IOV on 1GbE ports.
- If SR-IOV is disabled in BIOS or the Boot Manager, enabling SR-IOV from Intel PROSet will require a system reboot.
- Due to chipset limitations, not all systems or slots support SR-IOV. Below is a chart summarizing SR-IOV support on Dell server platforms.
- When an adapter is running in NPar mode, SR-IOV is limited to the root partition of each port.

- When an adapter is running in NPar mode, virtualization (SR-IOV) settings apply to all ports on the adapter, and to all partitions on each port. Changes made the virtualization settings on one port are applied to all ports on the adapter.

SR-IOV Support on Network Adapters

NDC, LOM, or Adapter	40Gbe	10Gbe	1Gbe
Intel® Ethernet Converged Network Adapter XL710-Q2	Yes	Yes	
Intel® Ethernet 40G 2P XL710 QSFP+ rNDC	Yes	Yes	
Intel® Ethernet 10G 4P X710-k bNDC		Yes	
Intel® Ethernet 10G 2P X710-k bNDC		Yes	
Intel® Ethernet 10G X710-k bNDC		Yes	
Intel® Converged Network Adapter X710		Yes	
Intel® Ethernet 10G 4P X710/I350 rNDC		Yes	No
Intel® Ethernet 10G 4P X710 SFP+ rNDC		Yes	
Intel® Ethernet 10G X710 rNDC		Yes	No
Intel® Ethernet 10G 2P X550-t Adapter		Yes	
Intel® Ethernet 10G 2P X540-t Adapter		Yes	
Intel Ethernet X540 DP 10Gb BT + I350 1Gb BT DP Network Daughter Card		Yes	No
Intel® Ethernet 10G 2P X520-k bNDC		Yes	
Intel® Ethernet 10G 2P X520 Adapter		Yes	
Intel® Ethernet X520 10GbE Dual Port KX4-KR Mezz		Yes	
Intel® Gigabit 4P I350-t rNDC		No	Yes
Intel® Gigabit 4P I350 bNDC		No	Yes
Intel® Gigabit 4P I350-t Mezz			Yes
Intel® Gigabit 2P I350-t Adapter			Yes
Intel® Gigabit 4P I350-t Adapter			Yes
PowerEdge C4130 LOMs			No
PowerEdge C6320 LOMs		Yes	
PowerEdge T620 LOMs			No
PowerEdge T630 LOMs			No
PowerEdge FC430 LOMs		No	Yes
PowerEdge R530XD LOMs			No

		Rack NDC		PCI Express Slot									
Dell Platform		**10 GbE Adapter**	**1 GbE Adapter**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
C4130				yes	yes								
C6320				yes									
R230				no	no								
R320				no	yes								
R330				no	no								
R420	1 x CPU			no	yes								
	2 x CPU			yes	yes								
R430				yes	yes								
R520	1 x CPU			no	yes	yes	yes						
	2 x CPU			yes	yes	yes	yes						
R530				yes	yes	yes	no	no					
R530XD				yes	yes	no							
R620				yes	yes	yes							
R630				yes	yes	yes							
R720XD		yes	no	yes	yes	yes	yes	yes	yes				
R720		yes	no	yes	yes	yes	yes	yes	yes	yes			
R730				yes	yes	yes	yes	yes	yes	yes			
R730XD				yes	yes	yes	yes	yes	yes				
R820		yes	no	yes	yes	yes	yes	yes	yes	yes			
R830				yes	yes	yes	yes	yes	yes				
R920		yes	no	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
R930				yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
T130				no	no	no	no						
T320				no	no	yes	yes		yes				
T330				no	no	no	no						
T420				no	no	yes	yes	yes	yes				
T430				no	no	yes	yes	yes	yes				
T620				yes	yes	no	yes	yes	yes	yes			
T630				yes	no	yes	yes	yes	yes	yes			

	Blade NDC		Mezzanine Slot	
Dell Platform	**10 GbE Adapter**	**1 GbE Adapter**	**B**	**C**
FC430		yes	yes	yes
FC630	yes		yes	yes

	Blade NDC		Mezzanine Slot	
Dell Platform	**10 GbE Adapter**	**1 GbE Adapter**	**B**	**C**
FC830	yes		yes	yes
M420	yes		yes	yes
M520	no		yes	yes
M620	yes		yes	yes
M630	yes		yes	yes
M630 for VRTX	yes			
M820	yes		yes	yes
M830	yes		yes	yes
M830 for VRTX	yes			

Supported platforms or slots are indicated by "yes." Unsupported are indicated by "no". Not applicable are indicated by blank cells.

Linux* Drivers for Intel® Ethernet Adapters

Overview

This release includes Linux Base Drivers for Intel® Network Connections. Specific information on building and installation, configuration, and command line parameters for these drivers are located in the following sections:

- igb Linux Driver for Intel® Gigabit Ethernet Adapters based on the 82575, 82576, I350, and I354 controllers
- ixgbe Linux Driver for Intel® 10 Gigabit Ethernet Adapters based on the 82598, 82599, and X540 controllers
- i40e Linux Driver for Intel® 10 Gigabit Ethernet Adapters based on the X710 and XL710 controllers

See the Supported Adapters section below to determine which driver to use.

These drivers are only supported as a loadable module. Intel is not supplying patches against the kernel source to allow for static linking of the driver. For questions related to hardware requirements, refer to System Requirements. All hardware requirements listed apply to use with Linux.

This release also includes support for Single Root I/O Virtualization (SR-IOV) drivers. More detail on SR-IOV can be found here. Intel recommends test-mode environments until industry hypervisors release production level support. The following drivers support the listed virtual function devices that can only be activated on kernels that support SR-IOV. SR-IOV requires the correct platform and OS support.

- igbvf Linux Driver for the Intel® Gigabit Family of Adapters for 82575, 82576, I350, and I354-based Gigabit Family of Adapters
- ixgbevf Linux Driver for the Intel® 10 Gigabit Family of Adapters for 82599 and X540-based 10 Gigabit Family of Adapters.
- i40e Linux Driver for the Intel® 10 Gigabit Family of Adapters for X710-based 10 Gigabit Family of Adapters and XL710-based 40 Gigabit Family of Adapters.

NOTES:

- On systems running Linux, Solaris, or ESXi, the base driver must be loaded for the Dell FW DUP to function correctly.
- The i40e driver does not support SR-IOV on ESXi 5.1.

Supported Adapters

The following Intel network adapters are compatible with the drivers in this release:

Devices supported by the igb Linux Base Driver

- Intel® Gigabit 4P I350-t rNDC
- Intel® Gigabit 4P X540/I350 rNDC
- Intel® Gigabit 4P X520/I350 rNDC
- Intel® Gigabit 4P I350-t Mezz
- Intel® Gigabit 4P X710/I350 rNDC
- Intel® Gigabit 4P I350-t bNDC
- Intel® Gigabit 2P I350-t Adapter
- Intel® Gigabit 4P I350-t Adapter
- Intel® Gigabit 4P I350 bNDC
- Intel® Ethernet Connection I354 1.0 GbE Backplane
- Intel® Gigabit 2P I350-t LOM
- Intel® Gigabit I350-t LOM
- Intel® Gigabit 2P I350 LOM

Devices Supported by the ixgbe Linux Base Driver

- Intel® Ethernet X520 10GbE Dual Port KX4-KR Mezz
- Intel® Ethernet 10G 2P X540-t Adapter
- Intel® Ethernet 10G 2P X550-t Adapter
- Intel® Ethernet 10G 4P X540/I350 rNDC
- Intel® Ethernet 10G 4P X520/I350 rNDC
- Intel® Ethernet 10G 2P X520-k bNDC

- Intel® Ethernet 10G 2P X520 Adapter
- Intel® Ethernet 10G X520 LOM

Devices Supported by the i40e Linux Base Driver

- Intel® Ethernet 10G 4P X710-k bNDC
- Intel® Ethernet 10G 2P X710-k bNDC
- Intel® Ethernet 10G X710-k bNDC
- Intel® Converged Network Adapter X710
- Intel® Ethernet 10G 4P X710/l350 rNDC
- Intel® Ethernet 10G 4P X710 SFP+ rNDC
- Intel® Ethernet 10G X710 rNDC
- Intel® Ethernet 40G 2P XL710 QSFP+ rNDC
- Intel® Ethernet Converged Network Adapter XL710-Q2

To verify your adapter is supported, find the board ID number on the adapter. Look for a label that has a barcode and a number in the format 123456-001 (six digits hyphen three digits). Match this to the list of numbers above.

For more information on how to identify your adapter or for the latest network drivers for Linux, see Customer Support.

Supported Linux Versions

Linux drivers are provided for the following distributions (only Intel® 64 versions are supported):

Red Hat Enterprise Linux (RHEL):

- Red Hat* Enterprise Linux* (RHEL) 7.2
- RHEL 7.1
- RHEL 6.7

SUSE Linux Enterprise Server (SLES):

- Novell* SUSE* Linux Enterprise Server (SLES) 12 SP1
- SLES 11 SP4

NIC Partitioning

Network Interface Card (NIC) Partitioning (NPar) allows network administrators to create multiple partitions for each physical port on a network adapter card, and to set different bandwidth allocations on each partition. To the network and operating system, each partition appears as a separate physical port on the adapter. This facilitates the reduction of switch port count and cabling complexity while maintaining network segmentation and isolation. In addition, flexible bandwidth allocation per partition allows for efficient use of the link.

NPar is supported by the following adapters, in combination with the Linux* i40e driver, and supports a maximum of 8 partitions on the following NICs:

- Intel® Ethernet 10G 4P X710-k bNDC
- Intel® Ethernet 10G 2P X710-k bNDC
- Intel® Ethernet 10G X710-k bNDC
- Intel® Converged Network Adapter X710
- Intel® Ethernet 10G 4P X710/l350 rNDC
- Intel® Ethernet 10G 4P X710 SFP+ rNDC
- Intel® Ethernet 10G X710 rNDC

NParEP Mode is a combination of NPar and PCIe ARI, and increases the maximum number of partitions on these adapters to 16 per NIC.

Both NPar and NParEP support SR-IOV on the adapters listed above (but only on the first partition).

NPar mode, and the NParEP extension, is activated through Device Settings in the System Setup Menu. The process of setting up port partitions is detailed in the Configuring NPar Mode section of the NPar chapter of this guide, here: Configuring NPar Mode.

Partition bandwidth allocation is set during the NPar Mode configuration pre-boot process described in the Configuring NPar Mode section. However, when NPar has been activated and the boot process has completed, the bandwidth allocations of each partition can be examined and/or reset via the Linux* command line. That procedure is described in the Linux Drivers section of the guide, which you can find here: Setting NPar Options In Linux*. Note however, that settings made through the Linux* command line are not persistent and will revert to the settings from the most recent system boot where bandwidth allocation settings were made should the system reboot.

Support

For general information and support, check with Customer Support.

If an issue is identified with the released source code on supported kernels with a supported adapter, email the specific information related to the issue to e1000-devel@lists.sf.net.

igb Linux* Driver for the Intel® Gigabit Adapters

igb Overview

This file describes the Linux* Base Driver for the Gigabit Intel® Network Connections based on the Intel® 82575EB, the Intel® 82576, the Intel® I350, and the Intel® I354. This driver supports the 2.6.x and 3.x kernels.

This driver is only supported as a loadable module. Intel is not supplying patches against the kernel source to allow for static linking of the driver. For questions related to hardware requirements, refer to System Requirements. All hardware requirements listed apply to use with Linux.

The following features are now available in supported kernels:

- Native VLANs
- Channel Bonding (teaming)
- SNMP

Adapter teaming is now implemented using the native Linux Channel bonding module. This is included in supported Linux kernels. Channel Bonding documentation can be found in the Linux kernel source: /documentation/networking/bonding.txt

The igb driver supports IEEE time stamping for kernels 2.6.30 and above.

Use ethtool, lspci, or ifconfig to obtain driver information. Instructions on updating the ethtool can be found in the Additional Configurations section later in this page.

igb Linux Base Driver Supported Devices

The following Intel network adapters are compatible with the igb driver in this release:

- Intel® Gigabit 4P I350-t rNDC
- Intel® Gigabit 4P X540/I350 rNDC
- Intel® Gigabit 4P X520/I350 rNDC
- Intel® Gigabit 4P I350-t Mezz
- Intel® Gigabit 4P X710/I350 rNDC
- Intel® Gigabit 4P I350-t bNDC
- Intel® Gigabit 2P I350-t Adapter
- Intel® Gigabit 4P I350-t Adapter
- Intel® Gigabit 4P I350 bNDC
- Intel® Ethernet Connection I354 1.0 GbE Backplane
- Intel® Gigabit 2P I350-t LOM
- Intel® Gigabit I350-t LOM
- Intel® Gigabit 2P I350 LOM

For information on how to identify your adapter or for the latest network drivers for Linux, see Customer Support.

Building and Installation

There are three methods for installing the igb driver:

- Install from Source Code
- Install Using KMP RPM
- Install Using KMOD RPM

Install from Source Code

To build a binary RPM* package of this driver, run 'rpmbuild -tb <filename.tar.gz>'. Replace <filename.tar.gz> with the specific filename of the driver.

NOTE:

- For the build to work properly it is important that the currently running kernel MATCH the version and configuration of the installed kernel source. If you have just recompiled your kernel, reboot the system.
- RPM functionality has only been tested in Red Hat distributions.

1. Download the base driver tar file to the directory of your choice. For example, use '`/home/username/igb`' or '`/usr/local/src/igb`'.
2. Untar/unzip the archive, where <x.x.x> is the version number for the driver tar:

   ```
   tar zxf igb-<x.x.x>.tar.gz
   ```

3. Change to the driver src directory, where `<x.x.x>` is the version number for the driver tar:

   ```
   cd igb-<x.x.x>/src/
   ```

4. Compile the driver module:

   ```
   # make install
   ```

   The binary will be installed as:

   ```
   /lib/modules/<KERNEL VERSION>/kernel/drivers/net/igb/igb.ko
   ```

   The install locations listed above are the default locations. This might differ for various Linux distributions. For more information, see the ldistrib.txt file included in the driver tar.
5. Install the module using the modprobe command:

   ```
   modprobe igb
   ```

   For 2.6 based kernels, make sure that the older igb drivers are removed from the kernel, before loading the new module:

   ```
   rmmod igb.ko; modprobe igb
   ```

6. Assign an IP address to and activate the Ethernet interface by entering the following, where <x> is the interface number:

   ```
   ifconfig eth<x> <IP_address> up
   ```

7. Verify that the interface works. Enter the following, where <IP_address> is the IP address for another machine on the same subnet as the interface that is being tested:

   ```
   ping <IP_address>
   ```

NOTE: Some systems have trouble supporting MSI and/or MSI-X interrupts. If your system needs to disable this type of interrupt, the driver can be built and installed with the command:

```
#make CFLAGS_EXTRA=-DDISABLE_PCI_MSI install
```

Normally, the driver generates an interrupt every two seconds. If interrupts are not received in cat /proc/interrupts for the ethX e1000e device, then this workaround may be necessary.

To build igb driver with DCA

If your kernel supports DCA, the driver will build by default with DCA enabled.

Install Using KMP RPM

NOTE: KMP is only supported on RHEL 6 and SLES 11.

The KMP RPMs update existing igb RPMs currently installed on the system. These updates are provided by SuSE in the SLES release. If an RPM does not currently exist on the system, the KMP will not install.

The RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

intel-<component name>-<component version>.<arch type>.rpm

For example, intel-igb-1.3.8.6-1.x86_64.rpm: igb is the component name; 1.3.8.6-1 is the component version; and x86_64 is the architecture type.

KMP RPMs are provided for supported Linux distributions. The naming convention for the included KMP RPMs is:

intel-<component name>-kmp-<kernel type>-<component version>_<kernel version>.<arch type>.rpm

For example, intel-igb-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm: igb is the component name; default is the kernel type; 1.3.8.6 is the component version; 2.6.27.19_5-1 is the kernel version; and x86_64 is the architecture type.

To install the KMP RPM, type the following two commands:

```
rpm -i <rpm filename>
rpm -i <kmp rpm filename>
```

For example, to install the igb KMP RPM package, type the following:

```
rpm -i intel-igb-1.3.8.6-1.x86_64.rpm
rpm -i intel-igb-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm
```

Install Using KMOD RPM

The KMOD RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

```
kmod-<driver name>-<version>-1.<arch type>.rpm
```

For example, kmod-igb-2.3.4-1.x86_64.rpm:

- igb is the driver name
- 2.3.4 is the version
- x86_64 is the architecture type

To install the KMOD RPM, go to the directory of the RPM and type the following command:

```
rpm -i <rpm filename>
```

For example, to install the igb KMOD RPM package from RHEL 6.4, type the following:

```
rpm -i kmod-igb-2.3.4-1.x86_64.rpm
```

Command Line Parameters

If the driver is built as a module, the following optional parameters are used by entering them on the command line with the modprobe command using this syntax:

```
modprobe igb [<option>=<VAL1>,<VAL2>,...]
```

A value (<VAL#>) must be assigned to each network port in the system supported by this driver. The values are applied to each instance, in function order. For example:

```
modprobe igb InterruptThrottleRate=16000,16000
```

In this case, there are two network ports supported by igb in the system. The default value for each parameter is generally the recommended setting, unless otherwise noted.

The following table contains parameters and possible values for modprobe commands:

Parameter Name	Valid Range/Settings	Default	Description
InterruptThrottleRate	0, 1, 3, 100-100000 (0=off, 1=dynamic, 3=d-dynamic conservative)	3	The driver can limit the number of interrupts per second that the adapter will generate for incoming packets. It does this by writing a value to the adapter that is based on the maximum number of interrupts that the adapter will generate per second. Setting InterruptThrottleRate to a value greater or equal to 100 will program the adapter to send out a maximum of that many interrupts per second, even if more packets have come in. This reduces interrupt load on the system and can lower CPU utilization under heavy load, but will increase latency as packets are not processed as quickly. The default behavior of the driver previously assumed a static InterruptThrottleRate value of 8000, providing a good fallback value for all traffic types, but lacking in small packet performance and latency. The driver has two adaptive modes (setting 1 or 3) in which it dynamically adjusts the InterruptThrottleRate value based on the traffic that it receives. After determining the type of incoming traffic in the last timeframe, it will adjust the InterruptThrottleRate to an appropriate value for that traffic. The algorithm classifies the incoming traffic every interval into classes. Once the class is determined, the InterruptThrottleRate value is adjusted to suit that traffic type the best. There are three classes defined: "Bulk traffic", for large amounts of packets of normal size; "Low latency", for small amounts of traffic and/or a significant percentage of small packets; and "Lowest latency", for almost completely small packets or minimal traffic. In dynamic conservative mode, the InterruptThrottleRate value is set to 4000 for traffic that falls in class "Bulk traffic". If traffic falls in the "Low latency" or "Lowest latency" class, the InterruptThrottleRate is increased stepwise to 20000. This default mode is suitable for most applications. For situations where low latency is vital such as cluster or grid computing, the algorithm can reduce latency even more when InterruptThrottleRate is set to mode 1. In this mode, which operates the same as mode 3, the InterruptThrottleRate will be increased stepwise to 70000 for traffic in class "Lowest latency". Setting InterruptThrottleRate to 0 turns off any interrupt moderation and may improve small packet latency, but is generally not suitable for bulk throughput traffic. **NOTE:** InterruptThrottleRate takes precedence over the TxAbsIntDelay and RxAbsIntDelay parameters. In other words, minimizing the receive and/or transmit absolute delays does not force the controller to generate more interrupts that what the Interrupt Throttle Rate allows.

<table>
<tr><th>Parameter Name</th><th>Valid Range/Settings</th><th>Default</th><th>Description</th></tr>
<tr><td>LLIPort</td><td>0-65535</td><td>0 (disabled)</td><td>LLIPort configures the port for Low Latency Interrupts (LLI).

Low Latency Interrupts allow for immediate generation of an interrupt upon processing receive packets that match certain criteria as set by the parameters described below. LLI parameters are not enabled when Legacy interrupts are used. You must be using MSI or MSI-X (see cat /proc/interrupts) to successfully use LLI.

For example, using LLIPort=80 would cause the board to generate an immediate interrupt upon receipt of any packet sent to TCP port 80 on the local machine.

CAUTION: Enabling LLI can result in an excessive number of interrupts/second that may cause problems with the system and in some cases may cause a kernel panic.</td></tr>
<tr><td>LLIPush</td><td>0-1</td><td>0 (disabled)</td><td>LLIPush can be set to enabled or disabled (default). It is most effective in an environment with many small transactions.

NOTE: Enabling LLIPush may allow a denial of service attack.</td></tr>
<tr><td>LLISize</td><td>0-1500</td><td>0 (disabled)</td><td>LLISize causes an immediate interrupt if the board receives a packet smaller than the specified size.</td></tr>
<tr><td>IntMode</td><td>0-2</td><td>2</td><td>This allows load time control over the interrupt type registered by the driver. MSI-X is required for multiple queue support. Some kernels and combinations of kernel .config options will force a lower level of interrupt support. 'cat/proc/interrupts' will show different values for each type of interrupt.

0 = Legacy Interrupts.
1 = MSI Interrupts.
2 = MSI-X interrupts (default).</td></tr>
<tr><td>RSS</td><td>0-8</td><td>1</td><td>0 = Assign up to whichever is less between the number of CPUs or the number of queues.
X = Assign X queue, where X is less than or equal to the maximum number of queues. The driver allows maximum supported queue value. For example, I350-based adapters allow RSS=8, where 8 is the maximum queues allowed.

NOTE: For 82575-based adapters, the maximum number of queues is 4. For 82576-based and newer adapters, it is 8.

This parameter is also affected by the VMDQ parameter in that it will limit the queues more.

<table>
<tr><td>Model</td><td>0</td><td>1</td><td>2</td><td>3+</td></tr>
<tr><td>82575</td><td>4</td><td>4</td><td>3</td><td>1</td></tr>
<tr><td>82576</td><td>8</td><td>2</td><td>2</td><td>2</td></tr>
</table>
</td></tr>
</table>

Parameter Name	Valid Range/Settings	Default	Description
VMDQ	0-4 for 82575-based adapters 0-8 for 82576-based adapters	0	This supports enabling VMDq pools, which is needed to support SR-IOV. This parameter is forced to 1 or more if the max_vfs module parameter is used. In addition, the number of queues available for RSS is limited if this is set to 1 or greater. 0 = Disabled 1 = Sets the netdev as pool 0 2 or greater = Add additional queues. However, these are currently not used. **NOTE:** When either SR-IOV mode or VMDq mode is enabled, hardware VLAN filtering and VLAN tag stripping/insertion will remain enabled.

max_vfs	0-7	0	This parameter adds support for SR-IOV. It causes the driver to spawn up to max_vfs worth of virtual function.
			If the value is greater than 0, it will force the VMDQ parameter to equal 1 or more.
			NOTE: When either SR-IOV mode or VMDq mode is enabled, hardware VLAN filtering and VLAN tag stripping/insertion will remain enabled. Please remove the old VLAN filter before the new VLAN filter is added. For example,
			`ip link set eth0 vf 0 vlan 100`
			`// set vlan 100 for VF 0`
			`ip link set eth0 vf 0 vlan 0 //`
			`Delete vlan 100`
			`ip link set eth0 vf 0 vlan 200`
			`// set a new vlan 200 for VF 0`

QueuePairs	0-1	1	This option can be overridden to 1 if there are not sufficient interrupts available. This can occur if any combination of RSS, VMDQ and max_vfs results in more than 4 queues being used. 0 = When MSI-X is enabled, the TX and RX will attempt to occupy separate vectors. 1 = TX and RX are paired onto one interrupt vector (default).
Node	0-n, where n is the number of the NUMA node that should be used to allocate memory for this adapter port. -1, uses the driver default of allocating memory on whichever processor is running modprobe.	-1 (off)	The Node parameter allows you to choose which NUMA node you want to have the adapter allocate memory from. All driver structures, in-memory queues, and receive buffers will be allocated on the node specified. This parameter is only useful when interrupt affinity is specified, otherwise some portion of the time the interrupt could run on a different core than the memory is allocated on, causing slower memory access and impacting throughput, CPU, or both.

Parameter Name	Valid Range/Settings	Default	Description
EEE	0-1	1 (enabled)	This option allows for the ability of IEEE802.3az, Energy Efficient Ethernet (EEE), to be advertised to the link partner on parts supporting EEE. A link between two EEE-compliant devices will result in periodic bursts of data followed by periods where the link is in an idle state. This Low Power Idle (LPI) state is supported in both 1 Gbps and 100 Mbps link speeds. **NOTES:** • EEE support requires auto-negotiation. • EEE is disabled by default on all I350-based adapters.
DMAC	0, 250, 500, 1000, 2000, 3000, 4000, 5000, 6000, 7000, 8000, 9000, 10000	0 (disabled)	Enables or disables DMA Coalescing feature. Values are in microseconds and increase the internal DMA Coalescing feature's internal timer. Direct Memory Access (DMA) allows the network device to move packet data directly to the system's memory, reducing CPU utilization. However, the frequency and random intervals at which packets arrive do not allow the system to enter a lower power state. DMA Coalescing allows the adapter to collect packets before it initiates a DMA event. This may increase network latency but also increases the chances that the system will enter a lower power state. Turning on DMA Coalescing may save energy with kernel 2.6.32 and later. This will impart the greatest chance for your system to consume less power. DMA Coalescing is effective in helping potentially saving the platform power only when it is enabled across all active ports. InterruptThrottleRate (ITR) should be set to dynamic. When ITR=0, DMA Coalescing is automatically disabled. A white paper containing information on how to best configure your platform is available on the Intel website.
MDD	0-1	1 (enabled)	The Malicious Driver Detection (MDD) parameter is only relevant for I350 devices operating in SR-IOV mode. When this parameter is set, the driver detects malicious VF driver and disables its TX/RX queues until a VF driver reset occurs.

Additional Configurations

Configuring the Driver on Different Distributions

Configuring a network driver to load properly when the system is started is distribution dependent. Typically, the configuration process involves adding an alias line to /etc/modules.conf or /etc/modprobe.conf as well as editing other system start up scripts and/or configuration files. Many Linux distributions ship with tools to make these changes for you. To learn the proper way to configure a network device for your system, refer to your distribution documentation. If during this process you are asked for the driver or module name, the name for the Linux Base Driver for the Intel Gigabit Family of Adapters is igb.

As an example, if you install the igb driver for two Intel Gigabit adapters (eth0 and eth1) and set the speed and duplex to 10 Full and 100 Half, add the following to modules.conf:

```
alias eth0 igb
alias eth1 igb
options igb IntMode=2,1
```

Viewing Link Messages

Link messages will not be displayed to the console if the distribution is restricting system messages. In order to see network driver link messages on your console, set dmesg to eight by entering the following:

```
dmesg -n 8
```

NOTE: This setting is not saved across reboots.

Jumbo Frames

Jumbo Frames support is enabled by changing the MTU to a value larger than the default of 1500 bytes. Use the ifconfig command to increase the MTU size. For example:

```
ifconfig eth<x> mtu 9000 up
```

This setting is not saved across reboots. The setting change can be made permanent by adding `MTU = 9000` to the file `/etc/sysconfig/network-scripts/ifcfg-eth<x>`, in Red Hat distributions. Other distributions may store this setting in a different location.

NOTES:

- Using Jumbo Frames at 10 or 100 Mbps may result in poor performance or loss of link.
- To enable Jumbo Frames, increase the MTU size on the interface beyond 1500.
- The maximum Jumbo Frames size is 9234 bytes, with a corresponding MTU size of 9216 bytes.

ethtool

The driver utilizes the ethtool interface for driver configuration and diagnostics, as well as displaying statistical information. ethtool version 3 or later is required for this functionality, although we strongly recommend downloading the latest version at: http://ftp.kernel.org/pub/software/network/ethtool/.

Speed and Duplex Configuration

In the default mode, an Intel® Network Adapter using copper connections will attempt to auto-negotiate with its link partner to determine the best setting. If the adapter cannot establish link with the link partner using auto-negotiation, you may need to manually configure the adapter and link partner to identical settings to establish link and pass packets. This should only be needed when attempting to link with an older switch that does not support auto-negotiation or one that has been forced to a specific speed or duplex mode.

Your link partner must match the setting you choose. Fiber-based adapters operate only in full duplex, and only at their native speed.

Speed and Duplex are configured through the ethtool* utility. ethtool is included with all versions of Red Hat after Red Hat 6.2. For other Linux distributions, download and install ethtool from the following website: http://ftp.kernel.org/pub/software/network/ethtool/.

CAUTION: Only experienced network administrators should force speed and duplex manually. The settings at the switch must always match the adapter settings. Adapter performance may suffer or your adapter may not operate if you configure the adapter differently from your switch.

Enabling Wake on LAN*

Wake on LAN (WoL) is configured through the ethtool* utility. ethtool is included with all versions of Red Hat after Red Hat 7.2. For other Linux distributions, download and install ethtool from the following website: http://ftp.kernel.org/pub/software/network/ethtool/.

For instructions on enabling WoL with ethtool, refer to the website listed above.

WoL will be enabled on the system during the next shut down or reboot. For this driver version, in order to enable WoL, the igb driver must be loaded prior to shutting down or suspending the system.

NOTE: Wake on LAN is only supported on port A of multi-port devices.

Multiqueue

In this mode, a separate MSI-X vector is allocated for each queue and one for “other” interrupts such as link status change and errors. All interrupts are throttled via interrupt moderation. Interrupt moderation must be used to avoid interrupt storms while the driver is processing one interrupt. The moderation value should be at least as large as the expected time for the driver to process an interrupt. Multiqueue is off by default.

MSI-X support is required for Multiqueue. If MSI-X is not found, the system will fallback to MSI or to Legacy interrupts. This driver supports multiqueue in kernel versions 2.6.24 and greater and supports receive multiqueue on all kernels supporting MSI-X.

NOTES:

- Do not use MSI-X with the 2.6.19 or 2.6.20 kernels. It is recommended to use the 2.6.21 or later kernel.
- Some kernels require a reboot to switch between single queue mode and multiqueue modes or vice-versa.

Large Receive Offload (LRO)

Large Receive Offload (LRO) is a technique for increasing inbound throughput of high-bandwidth network connections by reducing CPU overhead. It works by aggregating multiple incoming packets from a single stream into a larger buffer before they are passed higher up the networking stack, thus reducing the number of packets that have to be processed. LRO combines multiple Ethernet frames into a single receive in the stack, thereby potentially decreasing CPU utilization for receives.

NOTE: LRO requires 2.6.22 or later kernel version.

IGB_LRO is a compile time flag. It can be enabled at compile time to add support for LRO from the driver. The flag is used by adding CFLAGS_EXTRA="-DIGB_LRO" to the make file when it is being compiled. For example:

```
# make CFLAGS_EXTRA="-DIGB_LRO" install
```

You can verify that the driver is using LRO by looking at these counters in ethtool:

- lro_aggregated - count of total packets that were combined
- lro_flushed - counts the number of packets flushed out of LRO
- lro_no_desc - counts the number of times an LRO descriptor was not available for the LRO packet

NOTE: IPv6 and UDP are not supported by LRO.

IEEE 1588 Precision Time Protocol (PTP) Hardware Clock (PHC)

Precision Time Protocol (PTP) is an implementation of the IEEE 1588 specification allowing network cards to synchronize their clocks over a PTP-enabled network. It works through a series of synchronization and delay notification transactions that allow a software daemon to implement a PID controller to synchronize the network card clocks.

NOTE: PTP requires a 3.0.0 or later kernel version with PTP support enabled in the kernel and a user-space software daemon.

IGB_PTP is a compile time flag. The user can enable it at compile time to add support for PTP from the driver. The flag is used by adding CFLAGS_EXTRA="-DIGB_PTP" to the make file when it's being compiled:

```
make CFLAGS_EXTRA="-DIGB_PTP" install
```

NOTE: The driver will fail to compile if your kernel does not support PTP.

You can verify that the driver is using PTP by looking at the system log to see whether a PHC was attempted to be registered or not. If you have a kernel and version of ethtool with PTP support, you can check the PTP support in the driver by executing:

```
ethtool -T ethX
```

MAC and VLAN anti-spoofing feature

When a malicious driver attempts to send a spoofed packet, it is dropped by the hardware and not transmitted. An interrupt is sent to the PF driver notifying it of the spoof attempt.

When a spoofed packet is detected the PF driver will send the following message to the system log (displayed by the "dmesg" command):

```
Spoof event(s) detected on VF(n)
```

Where n=the VF that attempted to do the spoofing.

Setting MAC Address, VLAN and Rate Limit Using IProute2 Tool

You can set a MAC address of a Virtual Function (VF), a default VLAN and the rate limit using the IProute2 tool. Download the latest version of the iproute2 tool from Sourceforge if your version does not have all the features you require.

Known Issues

Using the igb Driver on 2.4 or Older 2.6 Based Kernels

Due to limited support for PCI Express in 2.4 kernels and older 2.6 kernels, the igb driver may run into interrupt related problems on some systems, such as no link or hang when bringing up the device.

It is recommend to use the newer 2.6 based kernels, as these kernels correctly configure the PCI Express configuration space of the adapter and all intervening bridges. If you are required to use a 2.4 kernel, use a 2.4 kernel newer than 2.4.30. For 2.6 kernels, use the 2.6.21 kernel or newer.

Alternatively, on 2.6 kernels you may disable MSI support in the kernel by booting with the "pci=nomsi" option or permanently disable MSI support in your kernel by configuring your kernel with CONFIG_PCI_MSI unset.

Compiling the Driver

When trying to compile the driver by running make install, the following error may occur:

"Linux kernel source not configured - missing version.h"

To solve this issue, create the version.h file by going to the Linux Kernel source tree and entering:

```
# make include/linux/version.h
```

Performance Degradation with Jumbo Frames

Degradation in throughput performance may be observed in some Jumbo frames environments. If this is observed, increasing the application's socket buffer size and/or increasing the /proc/sys/net/ipv4/tcp_*mem entry values may help. For more details, see the specific application documentation and in the text file /usr/src/linux*/Documentation/networking/ip-sysctl.txt.

Jumbo Frames on Foundry BigIron 8000 switch

There is a known issue using Jumbo frames when connected to a Foundry BigIron 8000 switch. This is a 3rd party limitation. If you experience loss of packets, lower the MTU size.

Multiple Interfaces on Same Ethernet Broadcast Network

Due to the default ARP behavior on Linux, it is not possible to have one system on two IP networks in the same Ethernet broadcast domain (non-partitioned switch) behave as expected. All Ethernet interfaces will respond to IP traffic for any IP address assigned to the system. This results in unbalanced receive traffic.

If you have multiple interfaces in a server, turn on ARP filtering by entering:

```
echo 1 > /proc/sys/net/ipv4/conf/all/arp_filter
```

(this only works if your kernel's version is higher than 2.4.5).

NOTE: This setting is not saved across reboots. However this configuration change can be made permanent through one of the following methods:

- Add the following line to `/etc/sysctl.conf`:

```
net.ipv4.conf.all.arp_filter = 1
```

- Install the interfaces in separate broadcast domains (either in different switches or in a switch partitioned to VLANs).

Disable Rx Flow Control with ethtool

In order to disable receive flow control using ethtool, you must turn off auto-negotiation on the same command line. For example:

```
ethtool -A eth? autoneg off rx
```

Unplugging Network Cable While ethtool -p is Running

In kernel versions 2.5.50 and later (including 2.6 kernel), unplugging the network cable while ethtool -p is running will cause the system to become unresponsive to keyboard commands, except for control-alt-delete. Restarting the system appears to be the only remedy.

Detected Tx Unit Hang in Quad Port Adapters

In some cases, ports 3 and 4 do not pass traffic and report "Detected Tx Unit Hang" followed by "NETDEV WATCHDOG: ethX: transmit timed out" errors. Ports 1 and 2 do not show any errors and will pass traffic.

This issue may be resolved by updating to the latest kernel and BIOS. You should use an OS that fully supports Message Signaled Interrupts (MSI) and make sure that MSI is enabled in your system's BIOS.

Do Not Use LRO when Routing Packets

Due to a known general compatibility issue with LRO and routing, do not use LRO when routing packets.

MSI-X Issues with Kernels Between 2.6.19 and 2.6.21 (inclusive)

Kernel panics and instability may be observed on any MSI-X hardware if you use irqbalance with kernels between 2.6.19 and 2.6.21. If these types of problems are encountered, you may disable the irqbalance daemon or upgrade to a newer kernel.

Rx Page Allocation Errors

Page allocation failure order:0 errors may occur under stress with kernels 2.6.25 and above. This is caused by the way the Linux kernel reports this stressed condition.

Host May Reboot after Removing PF when VF is Active in Guest

Using kernel versions earlier than 3.2, do not unload the PF driver with active VFs. Doing this will cause your VFs to stop working until you reload the PF driver and may cause a spontaneous reboot of your system.

igbvf Linux* Driver for the Intel® Gigabit Adapters

igbvf Overview

This driver supports upstream kernel versions 2.6.30 (or higher) x86_64.

The igbvf driver supports 82576-based and I350-based virtual function devices that can only be activated on kernels that support SR-IOV. SR-IOV requires the correct platform and OS support.

The igbvf driver requires the igb driver, version 2.0 or later. The igbvf driver supports virtual functions generated by the igb driver with a max_vfs value of 1 or greater. For more information on the max_vfs parameter refer to the section on the igb driver.

The guest OS loading the igbvf driver must support MSI-X interrupts.

This driver is only supported as a loadable module at this time. Intel is not supplying patches against the kernel source to allow for static linking of the driver. For questions related to hardware requirements, refer to the documentation supplied with your Intel Gigabit adapter. All hardware requirements listed apply to use with Linux.

Instructions on updating ethtool can be found in the section Additional Configurations later in this document.

NOTE: For VLANs, There is a limit of a total of 32 shared VLANs to 1 or more VFs.

igbvf Linux Base Driver Supported Devices

The following Intel network adapters are compatible with the igbvf driver in this release:

- Intel® Gigabit 4P I350-t rNDC
- Intel® Gigabit 4P X540/I350 rNDC
- Intel® Gigabit 4P X520/I350 rNDC
- Intel® Gigabit 4P I350-t Mezz
- Intel® Gigabit 4P X710/I350 rNDC
- Intel® Gigabit 4P I350-t bNDC
- Intel® Gigabit 2P I350-t Adapter
- Intel® Gigabit 4P I350-t Adapter
- Intel® Gigabit 4P I350 bNDC
- Intel® Gigabit 2P I350-t LOM
- Intel® Gigabit I350-t LOM
- Intel® Gigabit 2P I350 LOM

For information on how to identify your adapter or for the latest network drivers for Linux, see Customer Support.

Building and Installation

There are two methods for installing the igbvf driver:

- Install from Source Code
- Install Using KMP RPM

Install from Source Code

To build a binary RPM* package of this driver, run 'rpmbuild -tb <filename.tar.gz>'. Replace <filename.tar.gz> with the specific filename of the driver.

NOTE: For the build to work properly it is important that the currently running kernel MATCH the version and configuration of the installed kernel source. If you have just recompiled your kernel, reboot the system.

1. Download the base driver tar file to the directory of your choice. For example, use '`/home/username/igbvf`' or '`/usr/local/src/igbvf`'.
2. Untar/unzip the archive, where <x.x.x> is the version number for the driver tar:

   ```
   tar zxf igbvf-<x.x.x>.tar.gz
   ```

3. Change to the driver src directory, where `<x.x.x>` is the version number for the driver tar:

   ```
   cd igbvf-<x.x.x>/src/
   ```

4. Compile the driver module:

   ```
   # make install
   ```

   The binary will be installed as:

   ```
   /lib/modules/<KERNEL VERSION>/kernel/drivers/net/igbvf/igbvf.ko
   ```

   The install locations listed above are the default locations. This might differ for various Linux distributions. For more information, see the ldistrib.txt file included in the driver tar.
5. Install the module using the modprobe command:

   ```
   modprobe igbvf
   ```

   For 2.6 based kernels, make sure that the older igbvf drivers are removed from the kernel, before loading the new module:

   ```
   rmmod igbvf.ko; modprobe igbvf
   ```

6. Assign an IP address to and activate the Ethernet interface by entering the following, where <x> is the interface number:

   ```
   ifconfig eth<x> <IP_address> up
   ```

7. Verify that the interface works. Enter the following, where <IP_address> is the IP address for another machine on the same subnet as the interface that is being tested:

   ```
   ping <IP_address>
   ```

NOTE: Some systems have trouble supporting MSI and/or MSI-X interrupts. If your system needs to disable this type of interrupt, the driver can be built and installed with the command:

```
#make CFLAGS_EXTRA=-DDISABLE_PCI_MSI install
```

Normally, the driver generates an interrupt every two seconds. If interrupts are not received in cat /proc/interrupts for the ethX e1000e device, then this workaround may be necessary.

To build igbvf driver with DCA

If your kernel supports DCA, the driver will build by default with DCA enabled.

Install Using KMP RPM

NOTE: KMP is only supported on SLES11.

The KMP RPMs update existing igbvf RPMs currently installed on the system. These updates are provided by SuSE in the SLES release. If an RPM does not currently exist on the system, the KMP will not install.

The RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

intel-<component name>-<component version>.<arch type>.rpm

For example, intel-igbvf-1.3.8.6-1.x86_64.rpm: igbvf is the component name; 1.3.8.6-1 is the component version; and x86_64 is the architecture type.

KMP RPMs are provided for supported Linux distributions. The naming convention for the included KMP RPMs is:

intel-<component name>-kmp-<kernel type>-<component version>_<kernel version>.<arch type>.rpm

For example, intel-igbvf-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm: igbvf is the component name; default is the kernel type; 1.3.8.6 is the component version; 2.6.27.19_5-1 is the kernel version; and x86_64 is the architecture type.

To install the KMP RPM, type the following two commands:

```
rpm -i <rpm filename>
rpm -i <kmp rpm filename>
```

For example, to install the igbvf KMP RPM package, type the following:

```
rpm -i intel-igbvf-1.3.8.6-1.x86_64.rpm
rpm -i intel-igbvf-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm
```

Command Line Parameters

If the driver is built as a module, the following optional parameters are used by entering them on the command line with the modprobe command using this syntax:

```
modprobe igbvf [<option>=<VAL1>,<VAL2>,...]
```

A value (<VAL#>) must be assigned to each network port in the system supported by this driver. The values are applied to each instance, in function order. For example:

```
modprobe igbvf InterruptThrottleRate=16000,16000
```

In this case, there are two network ports supported by igb in the system. The default value for each parameter is generally the recommended setting, unless otherwise noted.

The following table contains parameters and possible values for modprobe commands:

Parameter Name	Valid Range/Settings	Default	Description
Interrupt-ThrottleRate	0, 1, 3, 100-100000 (0=off, 1=dynamic, 3=d-dynamic conservative)	3	The driver can limit the number of interrupts per second that the adapter will generate for incoming packets. It does this by writing a value to the adapter that is based on the maximum number of interrupts that the adapter will generate per second. Setting InterruptThrottleRate to a value greater or equal to 100 will program the adapter to send out a maximum of that many interrupts per second, even if more packets have come in. This reduces interrupt load on the system and can lower CPU utilization under heavy load, but will increase latency as packets are not processed as quickly.

Parameter Name	Valid Range/Settings	Default	Description
			The driver has two adaptive modes (setting 1 or 3) in which it dynamically adjusts the InterruptThrottleRate value based on the traffic that it receives. After determining the type of incoming traffic in the last time frame, it will adjust the InterruptThrottleRate to an appropriate value for that traffic. The algorithm classifies the incoming traffic every interval into classes. Once the class is determined, the InterruptThrottleRate value is adjusted to suit that traffic type the best. There are three classes defined: "Bulk traffic", for large amounts of packets of normal size; "Low latency", for small amounts of traffic and/or a significant percentage of small packets; and "Lowest latency", for almost completely small packets or minimal traffic. In dynamic conservative mode, the InterruptThrottleRate value is set to 4000 for traffic that falls in class "Bulk traffic". If traffic falls in the "Low latency" or "Lowest latency" class, the InterruptThrottleRate is increased stepwise to 20000. This default mode is suitable for most applications. For situations where low latency is vital such as cluster or grid computing, the algorithm can reduce latency even more when InterruptThrottleRate is set to mode 1. In this mode, which operates the same as mode 3, the InterruptThrottleRate will be increased stepwise to 70000 for traffic in class "Lowest latency". Setting InterruptThrottleRate to 0 turns off any interrupt moderation and may improve small packet latency, but is generally not suitable for bulk throughput traffic. **NOTES:** • Dynamic interrupt throttling is only applicable to adapters operating in MSI or Legacy interrupt mode, using a single receive queue. • When igbvf is loaded with default settings and multiple adapters are in use simultaneously, the CPU utilization may increase non-linearly. In order to limit the CPU utilization without impacting the overall throughput, it is recommended to load the driver as follows: `modprobe igbvf InterruptThrottleRate=3000,3000,3000` This sets the InterruptThrottleRate to 3000 interrupts/sec for the first, second, and third instances of the driver. The range of 2000 to 3000 interrupts per second works on a majority of systems and is a good starting point, but the optimal value will be platform-specific. If CPU utilization is not a concern, use default driver settings.

Additional Configurations

Configuring the Driver on Different Distributions

Configuring a network driver to load properly when the system is started is distribution dependent. Typically, the configuration process involves adding an alias line to /etc/modules.conf or /etc/modprobe.conf as well as editing other system startup scripts and/or configuration files. Many Linux distributions ship with tools to make these changes for you. To learn the proper way to configure a network device for your system, refer to your distribution documentation. If during this process you are asked for the driver or module name, the name for the Linux Base Driver for the Intel Gigabit Family of Adapters is igbvf.

As an example, if you install the igbvf driver for two Intel Gigabit adapters (eth0 and eth1) and want to set the interrupt mode to MSI-X and MSI, respectively, add the following to modules.conf or /etc/modprobe.conf:

```
alias eth0 igbvf
alias eth1 igbvf
options igbvf InterruptThrottleRate=3,1
```

Viewing Link Messages

Link messages will not be displayed to the console if the distribution is restricting system messages. In order to see network driver link messages on your console, set dmesg to eight by entering the following:

```
dmesg -n 8
```

NOTE: This setting is not saved across reboots.

Jumbo Frames

Jumbo Frames support is enabled by changing the MTU to a value larger than the default of 1500 bytes. Use the ifconfig command to increase the MTU size. For example:

```
ifconfig eth<x> mtu 9000 up
```

This setting is not saved across reboots. The setting change can be made permanent by adding `MTU = 9000` to the file `/etc/sysconfig/network-scripts/ifcfg-eth<x>`, in Red Hat distributions. Other distributions may store this setting in a different location.

NOTES:

- Using Jumbo Frames at 10 or 100 Mbps may result in poor performance or loss of link.
- To enable Jumbo Frames, increase the MTU size on the interface beyond 1500.
- The maximum Jumbo Frames size is 9234 bytes, with a corresponding MTU size of 9216 bytes.

ethtool

The driver utilizes the ethtool interface for driver configuration and diagnostics, as well as displaying statistical information. ethtool version 3 or later is required for this functionality, although we strongly recommend downloading the latest version at: http://ftp.kernel.org/pub/software/network/ethtool/.

Known Issues

Compiling the Driver

When trying to compile the driver by running make install, the following error may occur:

"Linux kernel source not configured - missing version.h"

To solve this issue, create the version.h file by going to the Linux Kernel source tree and entering:

```
# make include/linux/version.h
```

Multiple Interfaces on Same Ethernet Broadcast Network

Due to the default ARP behavior on Linux, it is not possible to have one system on two IP networks in the same Ethernet broadcast domain (non-partitioned switch) behave as expected. All Ethernet interfaces will respond to IP traffic for any IP address assigned to the system. This results in unbalanced receive traffic.

If you have multiple interfaces in a server, turn on ARP filtering by entering:

```
echo 1 > /proc/sys/net/ipv4/conf/all/arp_filter
```

(this only works if your kernel's version is higher than 2.4.5).

NOTE: This setting is not saved across reboots. However this configuration change can be made permanent through one of the following methods:

- Add the following line to `/etc/sysctl.conf`:

  ```
  net.ipv4.conf.all.arp_filter = 1
  ```

- Install the interfaces in separate broadcast domains (either in different switches or in a switch partitioned to VLANs).

Do Not Use LRO when Routing Packets

Due to a known general compatibility issue with LRO and routing, do not use LRO when routing packets.

MSI-X Issues with Kernels Between 2.6.19 and 2.6.21 (inclusive)

Kernel panics and instability may be observed on any MSI-X hardware if you use irqbalance with kernels between 2.6.19 and 2.6.21. If these types of problems are encountered, you may disable the irqbalance daemon or upgrade to a newer kernel.

Rx Page Allocation Errors

Page allocation failure order:0 errors may occur under stress with kernels 2.6.25 and above. This is caused by the way the Linux kernel reports this stressed condition.

Unloading Physical Function (PF) Driver Causes System Reboots when VM is Running and VF is Loaded on the VM

Do not unload the PF driver (igb) while VFs are assigned to guests.

Host May Reboot after Removing PF when VF is Active in Guest

Using kernel versions earlier than 3.2, do not unload the PF driver with active VFs. Doing this will cause your VFs to stop working until you reload the PF driver and may cause a spontaneous reboot of your system.

ixgbe Linux* Driver for the Intel® 10 Gigabit Server Adapters

ixgbe Overview

WARNING: By default, the ixgbe driver complies with the Large Receive Offload (LRO) feature enabled. This option offers the lowest CPU utilization for receives but is incompatible with routing/ip forwarding and bridging. If enabling ip forwarding or bridging is a requirement, it is necessary to disable LRO using compile time options as noted in the LRO section later in this section. The result of not disabling LRO when combined with ip forwarding or bridging can be low throughput or even a kernel panic.

This file describes the Linux* Base Driver for the 10 Gigabit Intel® Network Connections. This driver supports the 2.6.x kernels and includes support for any Linux supported system, including X86_64, i686 and PPC.

This driver is only supported as a loadable module. Intel is not supplying patches against the kernel source to allow for static linking of the driver. A version of the driver may already be included by your distribution or the kernel. For questions related to hardware requirements, refer to System Requirements. All hardware requirements listed apply to use with Linux.

The following features are now available in supported kernels:

- Native VLANs
- Channel Bonding (teaming)
- SNMP
- Generic Receive Offload
- Data Center Bridging

Adapter teaming is now implemented using the native Linux Channel bonding module. This is included in supported Linux kernels. Channel Bonding documentation can be found in the Linux kernel source: /documentation/networking/bonding.txt

Use ethtool, lspci, or ifconfig to obtain driver information. Instructions on updating the ethtool can be found in the Additional Configurations section later in this page.

ixgbe Linux Base Driver Supported Devices

The following Intel network adapters are compatible with the Linux driver in this release:

- Intel® Ethernet X520 10GbE Dual Port KX4-KR Mezz
- Intel® Ethernet 10G 2P X540-t Adapter
- Intel® Ethernet 10G 2P X550-t Adapter
- Intel® Ethernet 10G 4P X540/I350 rNDC
- Intel® Ethernet 10G 4P X520/I350 rNDC
- Intel® Ethernet 10G 2P X520-k bNDC
- Intel® Ethernet 10G 2P X520 Adapter
- Intel® Ethernet 10G X520 LOM

For information on how to identify your adapter or for the latest network drivers for Linux, see Customer Support.

SFP+ Devices with Pluggable Optics

NOTE: For 92500-based SFP+ fiber adapters, using "`ifconfig down`" turns off the laser. "`ifconfig up`" turns the laser on.

For information on using SFP+ devices with pluggable optics, click here.

Building and Installation

There are three methods for installing the Linux driver:

- Install from Source Code
- Install Using KMP RPM
- Install Using KMOD RPM

Install from Source Code

To build a binary RPM* package of this driver, run 'rpmbuild -tb <filename.tar.gz>'. Replace <filename.tar.gz> with the specific filename of the driver.

NOTES:

- For the build to work properly it is important that the currently running kernel MATCH the version and configuration of the installed kernel source. If you have just recompiled your kernel, reboot the system.
- RPM functionality has only been tested in Red Hat distributions.

1. Download the base driver tar file to the directory of your choice. For example, use '`/home/username/ixgbe`' or '`/usr/local/src/ixgbe`'.
2. Untar/unzip the archive, where <x.x.x> is the version number for the driver tar:

```
tar zxf ixgbe-<x.x.x>.tar.gz
```

3. Change to the driver src directory, where <x.x.x> is the version number for the driver tar:

   ```
   cd ixgbe-<x.x.x>/src/
   ```

4. Compile the driver module:

   ```
   make install
   ```

   The binary will be installed as: /lib/modules/<KERNEL VERSION>/kernel/drivers/net/ixgbe/ixgbe.ko

   The install locations listed above are the default locations. This might differ for various Linux distributions. For more information, see the ldistrib.txt file included in the driver tar.

   NOTE: IXGBE_NO_LRO is a compile time flag. The user can enable it at compile time to remove support for LRO from the driver. The flag is used by adding `CFLAGS_EXTRA=-"DIXGBE_NO_LRO"` to the make file when it is being compiled. For example:

   ```
   make CFLAGS_EXTRA="-DIXGBE_NO_LRO" install
   ```

5. Install the module using the modprobe command for kernel 2.6.x:

   ```
   modprobe ixgbe <parameter>=<value>
   ```

   For 2.6 based kernels, make sure that the older ixgbe drivers are removed from the kernel, before loading the new module:

   ```
   rmmod ixgbe; modprobe ixgbe
   ```

6. Assign an IP address to and activate the Ethernet interface by entering the following, where <x> is the interface number:

   ```
   ifconfig eth<x> <IP_address> netmask <netmask>
   ```

7. Verify that the interface works. Enter the following, where <IP_address> is the IP address for another machine on the same subnet as the interface that is being tested:

   ```
   ping <IP_address>
   ```

Install Using KMP RPM

NOTE: KMP is only supported on RHEL 6 and SLES11.

The KMP RPMs update existing ixgbe RPMs currently installed on the system. These updates are provided by SuSE in the SLES release. If an RPM does not currently exist on the system, the KMP will not install.

The RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

intel-<component name>-<component version>.<arch type>.rpm

For example, intel-ixgbe-1.3.8.6-1.x86_64.rpm: ixgbe is the component name; 1.3.8.6-1 is the component version; and x86_64 is the architecture type.

KMP RPMs are provided for supported Linux distributions. The naming convention for the included KMP RPMs is:

intel-<component name>-kmp-<kernel type>-<component version>_<kernel version>.<arch type>.rpm

For example, intel-ixgbe-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm: ixgbe is the component name; default is the kernel type; 1.3.8.6 is the component version; 2.6.27.19_5-1 is the kernel version; and x86_64 is the architecture type.

To install the KMP RPM, type the following two commands:

```
rpm -i <rpm filename>
rpm -i <kmp rpm filename>
```

For example, to install the ixgbe KMP RPM package, type the following:

```
rpm -i intel-ixgbe-1.3.8.6-1.x86_64.rpm
rpm -i intel-ixgbe-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm
```

Install Using KMOD RPM

The KMOD RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

```
kmod-<driver name>-<version>-1.<arch type>.rpm
```

For example, kmod-ixgbe-2.3.4-1.x86_64.rpm:

- ixgbe is the driver name
- 2.3.4 is the version
- x86_64 is the architecture type

To install the KMOD RPM, go to the directory of the RPM and type the following command:

```
rpm -i <rpm filename>
```

For example, to install the ixgbe KMOD RPM package from RHEL 6.4, type the following:

```
rpm -i kmod-ixgbe-2.3.4-1.x86_64.rpm
```

Command Line Parameters

If the driver is built as a module, the following optional parameters are used by entering them on the command line with the modprobe command using this syntax:

```
 modprobe ixgbe [<option>=<VAL1>,<VAL2>,...]
```

For example:

```
modprobe ixgbe InterruptThrottleRate=16000,16000
```

The default value for each parameter is generally the recommended setting, unless otherwise noted.

The following table contains parameters and possible values for modprobe commands:

Parameter Name	Valid Range/Settings	Default	Description
RSS	0 - 16	1	Receive Side Scaling allows multiple queues for receiving data. 0 = Sets the descriptor queue count to the lower value of either the number of CPUs or 16. 1 - 16 = Sets the descriptor queue count to 1 - 16. RSS also effects the number of transmit queues allocated on 2.6.23 and newer kernels with CONFIG_NET_MULTIQUEUE set in the kernel .config file. CONFIG_NETDEVICES_ MULTIQUEUE is only supported in kernels 2.6.23 to 2.6.26. For kernels 2.6.27 or newer, other options enable multiqueue. **NOTE:** The RSS parameter has no effect on 82599-based adapters unless the FdirMode parameter is simultaneously used to disable Flow Director. See Intel® Ethernet Flow Director section for more detail.
MQ	0, 1	1	Multi Queue support. 0 = Disables Multiple Queue support. 1 = Enables Multiple Queue support (prerequisite for RSS).
IntMode	0 - 2	2	Interrupt mode controls the allowed load time control over the type of interrupt registered for by the driver. MSI-X is required for multiple queue support, and some kernels and combinations of kernel .config options will force a lower level of interrupt support. 'cat /proc/interrupts' will show different values for each type of interrupt. 0 = Legacy interrupt 1 = MSI 2 = MSIX
InterruptThrottleRate	956 - 488,281 (0=off, 1=d-dynamic)	1	Interrupt Throttle Rate (interrupts/second). The ITR parameter controls how many interrupts each interrupt vector can generate per second. Increasing ITR lowers latency at the cost of increased CPU utilization, though it may help throughput in some circumstances. 0 = This turns off any interrupt moderation and may improve small packet latency. However, it is generally not suitable for bulk throughput traffic due to the increased CPU utilization of the higher interrupt rate. **NOTES:** • For 82599-based adapters, disabling InterruptThrottleRate will also result in the driver disabling HW RSC. • For 82598-based adapters, disabling InterruptThrottleRate will also result in disabling LRO. 1 = Dynamic mode attempts to moderate interrupts per vector while maintaining very low latency. This can sometimes cause extra CPU utilization when in dynamic mode. If planning on deploying ixgbe in a latency sensitive environment, please consider this parameter.

Parameter Name	Valid Range/Settings	Default	Description
LLI			Low Latency Interrupts allow for immediate generation of an interrupt upon processing receive packets that match certain criteria as set by the parameters described below. LLI parameters are not enabled when Legacy interrupts are used. You must be using MSI or MSI-X (see cat /proc/interrupts) to successfully use LLI.
LLIPort	0 - 65535	0 (disabled)	LLI is configured with the LLIPort command-line parameter, which specifies which TCP should generate Low Latency Interrupts. For example, using LLIPort=80 would cause the board to generate an immediate interrupt upon receipt of any packet sent to TCP port 80 on the local machine. **WARNING:** Enabling LLI can result in an excessive number of interrupts/second that may cause problems with the system and in some cases may cause a kernel panic.
LLIPush	0 - 1	0 (disabled)	LLIPush can be set to be enabled or disabled (default). It is most effective in an environment with many small transactions. **NOTE:** Enabling LLIPush may allow a denial of service attack.
LLISize	0 - 1500	0 (disabled)	LLISize causes an immediate interrupt if the board receives a packet smaller than the specified size.
LLIEType	0 - x8FFF	0 (disabled)	Low Latency Interrupt Ethernet Protocol Type.
LLIVLANP	0 - 7	0 (disabled)	Low Latency Interrupt on VLAN Priority Threshold.
Flow Control			Flow Control is enabled by default. If you want to disable a flow control capable link partner, use ethtool: `ethtool -A eth? autoneg off rx off tx off` **NOTE:** For 82598 backplane cards entering 1 Gbps mode, flow control default behavior is changed to off. Flow control in 1 Gbps mode on these devices can lead to transmit hangs.
Intel® Ethernet Flow Director			**NOTE:** Flow director parameters are only supported on kernel versions 2.6.30 or later. Flow control in 1 Gbps mode on these devices can lead to transmit hangs. This supports advanced filters that direct receive packets by their flows to different queues and enables tight control on routing a flow in the platform. It matches flows and CPU cores for flow affinity and supports multiple parameters for flexible flow classification and load balancing. The flow director is enabled only if the kernel is multiple TX queue capable. An included script (set_irq_affinity.sh) automates setting the IRQ to CPU affinity. To verify that the driver is using Flow Director, look at the counter in ethtool: fdir_miss and fdir_match.

Parameter Name	Valid Range/Settings	Default	Description

			Other ethtool Commands:
			To enable Flow Director
			`ethtool -K ethX ntuple on`
			To add a filter, use -U switch
			`ethtool -U ethX flow-type tcp4 src-ip 192.168.0.100 action 1`
			To see the list of filters currently present
			`ethtool -u ethX`
			Perfect Filter: Perfect filter is an interface to load the filter table that funnels all flow into queue_0 unless an alternative queue is specified using "action." In that case, any flow that matches the filter criteria will be directed to the appropriate queue.
			Support for Virtual Function (VF) is via the user-data field. You must update to the version of ethtool built for the 2.6.40 kernel. Perfect Filter is supported on all kernels 2.6.30 and later. Rules may be deleted from the table itself. This is done via "ethtool -U ethX delete N" where N is the rule number to be deleted.
			NOTE: Flow Director Perfect Filters can run in single queue mode, when SR-IOV is enabled, or when DCB is enabled.
			If the queue is defined as -1, the filter will drop matching packets.
			To account for filter matches and misses, there are two stats in ethtool: fdir_match and fdir_miss. In addition, rx_queue_N_ packets shows the number of packets processed by the Nth queue.
			NOTES:
			• Receive Packet Steering (RPS) and Receive Flow Steering (RFS) are not compatible with Flow Director. If Flow Director is enabled, these will be disabled.
			• For VLAN Masks only 4 masks are supported.
			• Once a rule is defined, you must supply the same fields and masks (if masks are specified).
			Support for UDP RSS
			This feature adds an ON/OFF switch for hashing over certain flow types. You can't turn on anything other than UDP. The default setting is disabled. We only support enabling/disabling hashing on ports for UDP over IPv4 (udp4) or IPv6 (udp6).
			NOTE: Fragmented packets may arrive out of order when RSS UDP support is configured.
			Supported ethtool Commands and Options
			`-n --show-nfc`
			Retrieves the receive network flow classification configurations.
			`rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6`
			Retrieves the hash options for the specified network traffic type.

Parameter Name	Valid Range/Settings	Default	Description

			`-N --config-nfc`
			Configures the receive network flow classification.
			`rx-flow-hash tcp4|ud-p4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6 m|v|t|s|d|f|n|r...`
			Configures the hash options for the specified network traffic type.
			`udp4`
			UDP over IPv4
			`udp6`
			UDP over IPv6
			`f`
			Hash on bytes 0 and 1 of the Layer 4 header of the rx packet.
			`n`
			Hash on bytes 2 and 3 of the Layer 4 header of the rx packet.
			The following is an example using udp4 (UDP over IPv4):
			To include UDP port numbers in RSS hashing run:
			`ethtool -N eth1 rx-flow-hash udp4 sdfn`
			To exclude UDP port numbers from RSS hashing run:
			`ethtool -N eth1 rx-flow-hash udp4 sd`
			To display UDP hashing current configuration run:
			`ethtool -n eth1 rx-flow-hash udp4`
			The results of running that call will be the following, if UDP hashing is enabled:
			UDP over IPv4 flows use these fields for computing Hash flow key:
			IP SA
			IP DA
			L4 bytes 0 & 1 [TCP/UDP src port]
			L4 bytes 2 & 3 [TCP/UDP dst port]
			The results if UDP hashing is disabled would be:
			UDP over IPv4 flows use these fields for computing Hash flow key:
			IP SA
			IP DA
			The following two parameters impact Flow Director: FdirPballoc and AtrSampleRate.

FdirPballoc	0 - 2	0 (64k)	Flow Allocated Packet Buffer Size. 0 = 64k 1 = 128k 2 = 256k
AtrSampleRate	1 - 100	20	Software ATR Tx Packet Sample Rate. For example, when set to 20, every 20th packet is sampled to determine if the packet will create a new flow.

Parameter Name	Valid Range/Settings	Default	Description

max_vfs	1 - 63	0	This parameter adds support for SR-IOV. It causes the driver to spawn up to max_vfs worth of virtual function.
			
			If the value is greater than 0, it will also force the VMDq parameter to be 1 or more.
			
			NOTE: When either SR-IOV mode or VMDq mode is enabled, hardware VLAN filtering and VLAN tag stripping/insertion will remain enabled. Please remove the old VLAN filter before the new VLAN filter is added. For example:
			
			`ip link set eth0 vf 0 vlan 100`
			`// set vlan 100 for VF 0`
			`ip link set eth0 vf 0 vlan 0 //`
			`Delete vlan 100`
			`ip link set eth0 vf 0 vlan 200`
			`// set a new vlan 200 for VF 0`
			
			The parameters for the driver are referenced by position. So, if you have a dual port 82599-based adapter and you want N virtual functions per port, you must specify a number for each port with each parameter separated by a comma.
			
			For example: `modprobe ixgbe max_vfs=63,63`
			
			NOTE: If both 82598 and 82599-based adapters are installed on the same machine, you must be careful in loading the driver with the parameters. Depending on system configuration, number of slots, etc., it is impossible to predict in all cases where the positions would be on the command line and the user will have to specify zero in those positions occupied by an 82598 port.
			
			With kernel 3.6, the driver supports the simultaneous usage of max_vfs and DCB features, subject to the constraints described below. Prior to kernel 3.6, the driver did not support the simultaneous operation of max_vfs > 0 and the DCB features (multiple traffic classes utilizing Priority Flow Control and Extended Transmission Selection).
			
			When DCB is enabled, network traffic is transmitted and received through multiple traffic classes (packet buffers in the NIC). The traffic is associated with a specific class based on priority, which has a value of 0 through 7 used in the VLAN tag. When SR-IOV is not enabled, each traffic class is associated with a set of RX/TX descriptor queue pairs. The number of queue pairs for a given traffic class depends on the hardware configuration. When SR-IOV is enabled, the descriptor queue pairs are grouped into pools. The Physical Function (PF) and each Virtual Function (VF) is allocated a pool of RX/TX descriptor queue pairs. When multiple traffic classes are configured (for example, DCB is enabled), each pool contains a queue pair from each traffic class. When a single traffic class is configured in the hardware, the pools contain multiple queue pairs from the single traffic class.
			
			The number of VFs that can be allocated depends on the number of traffic classes that can be enabled. The configurable number of traffic classes for each enabled VF is as follows:

Parameter Name	Valid Range/Settings	Default	Description
			0 - 15 VFs = Up to 8 traffic classes, depending on device support 16 - 31 VFs = Up to 4 traffic classes 32 - 63 = 1 traffic class When VFs are configured, the PF is allocated one pool as well. The PF supports the DCB features with the constraint that each traffic class will only use a single queue pair. When zero VFs are configured, the PF can support multiple queue pairs per traffic class.
VMDQ	1-16	1 (disabled)	This provides the option for turning VMDQ on or off. Values 2 through 16 enable VMDQ with the descriptor queues set to the specified value.
L2LBen	0-1	1 (enabled)	This parameter controls the internal switch (L2 loopback between pf and vf). By default the switch is enabled.

Additional Configurations

Configuring the Driver on Different Distributions

Configuring a network driver to load properly when the system is started is distribution dependent. Typically, the configuration process involves adding an alias line to /etc/modules.conf or /etc/modprobe.conf as well as editing other system startup scripts and/or configuration files. Many Linux distributions ship with tools to make these changes for you. To learn the proper way to configure a network device for your system, refer to your distribution documentation. If during this process you are asked for the driver or module name, the name for the Linux Base Driver for the Intel® 10 Gigabit PCI Express Family of Adapters is ixgbe.

Viewing Link Messages

Link messages will not be displayed to the console if the distribution is restricting system messages. In order to see network driver link messages on your console, set dmesg to eight by entering the following:

```
dmesg -n 8
```

NOTE: This setting is not saved across reboots.

Jumbo Frames

Jumbo Frames support is enabled by changing the MTU to a value larger than the default of 1500 bytes. The maximum value for the MTU is 9710. Use the ifconfig command to increase the MTU size. For example, enter the following where <x> is the interface number:

```
ifconfig ethx mtu 9000 up
```

This setting is not saved across reboots. The setting change can be made permanent by adding `MTU = 9000` to the file `/etc/sysconfig/network-scripts/ifcfg-eth<x>` for RHEL or to the file `/etc/sysconfig/network/<config_file>` for SLES.

The maximum MTU setting for Jumbo Frames is 9710. This value coincides with the maximum Jumbo Frames size of 9728. This driver will attempt to use multiple page sized buffers to receive each jumbo packet. This should help to avoid buffer starvation issues when allocating receive packets.

For 82599-based network connections, if you are enabling jumbo frames in a virtual function (VF), jumbo frames must first be enabled in the physical function (PF). The VF MTU setting cannot be larger than the PF MTU.

ethtool

The driver uses the ethtool interface for driver configuration and diagnostics, as well as displaying statistical information. The latest ethtool version is required for this functionality.

The latest release of ethtool can be found at: http://sourceforge.net/projects/gkernel.

NAPI

NAPI (Rx polling mode) is supported in the ixgbe driver.

See ftp://robur.slu.se/pub/Linux/net-development/NAPI/usenix-paper.tgz for more information on NAPI.

Large Receive Offload (LRO)

Large Receive Offload (LRO) is a technique for increasing inbound throughput of high-bandwidth network connections by reducing CPU overhead. It works by aggregating multiple incoming packets from a single stream into a larger buffer before they are passed higher up the networking stack, thus reducing the number of packets that have to be processed. LRO combines multiple Ethernet frames into a single receive in the stack, thereby potentially decreasing CPU utilization for receives.

IXGBE_NO_LRO is a compile time flag. The user can enable it at compile time to remove support for LRO from the driver. The flag is used by adding CFLAGS_EXTRA="-DIXGBE_NO_LRO" to the make file when it's being compiled.

```
make CFLAGS_EXTRA="-DIXGBE_NO_LRO" install
```

You can verify that the driver is using LRO by looking at these counters in ethtool:

- lro_flushed - the total number of receives using LRO.
- lro_coal - counts the total number of Ethernet packets that were combined.

HW RSC

82599-based adapters support hardware based receive side coalescing (RSC) which can merge multiple frames from the same IPv4 TCP/IP flow into a single structure that can span one or more descriptors. It works similarly to software large receive offload technique. By default HW RSC is enabled, and SW LRO can not be used for 82599-based adapters unless HW RSC is disabled.

IXGBE_NO_HW_RSC is a compile time flag that can be enabled at compile time to remove support for HW RSC from the driver. The flag is used by adding CFLAGS_EXTRA="-DIXGBE_NO_HW_RSC" to the make file when it is being compiled.

```
make CFLAGS_EXTRA="-DIXGBE_NO_HW_RSC" install
```

You can verify that the driver is using HW RSC by looking at the counter in ethtool:

```
hw_rsc_count - counts the total number of Ethernet packets that were being com-
bined.
```

rx_dropped_backlog

When in a non-Napi (or Interrupt) mode, this counter indicates that the stack is dropping packets. There is an adjustable parameter in the stack that allows you to adjust the amount of backlog. We recommend increasing the netdev_max_backlog if the counter goes up.

```
# sysctl -a |grep netdev_max_backlog

net.core.netdev_max_backlog = 1000

# sysctl -e net.core.netdev_max_backlog=10000

net.core.netdev_max_backlog = 10000
```

Flow Control

Flow control is disabled by default. To enable it, use ethtool:

```
ethtool -A eth? autoneg off rx on tx on
```

NOTE: You must have a flow control capable link partner.

MAC and VLAN Anti-spoofing Feature

When a malicious driver attempts to send a spoofed packet, it is dropped by the hardware and not transmitted. An interrupt is sent to the PF driver notifying it of the spoof attempt. When a spoofed packet is detected the PF driver will send the following message to the system log (displayed by the "dmesg" command):

```
ixgbe ethx: ixgbe_spoof_check: n spoofed packets detected
```

Where x=the PF interface# and n=the VF that attempted to do the spoofing.

NOTE: This feature can be disabled for a specific Virtual Function (VF).

Support for UDP RSS

This feature adds an ON/OFF switch for hashing over certain flow types. The default setting is disabled. NOTE: Fragmented packets may arrive out of order when RSS UDP support is configured.

Supported ethtool Commands and Options

`-n --show-nfc`
Retrieves the receive network flow classification configurations.

`rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6`
Retrieves the hash options for the specified network traffic type.

`-N --config-nfc`
Configures the receive network flow classification.

`rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6 m|v|t|s|d|f|n|r...`
Configures the hash options for the specified network traffic type.

udp4 UDP over IPv4
udp6 UDP over IPv6

f Hash on bytes 0 and 1 of the Layer 4 header of the rx packet.
n Hash on bytes 2 and 3 of the Layer 4 header of the rx packet.

Known Issues

Driver Compilation

When trying to compile the driver by running make install, the following error may occur: "Linux kernel source not configured - missing version.h"

To solve this issue, create the version.h file by going to the Linux source tree and entering:

```
make include/linux/version.h
```

Do Not Use LRO when Routing Packets

Due to a known general compatibility issue with LRO and routing, do not use LRO when routing packets.

Performance Degradation with Jumbo Frames

Degradation in throughput performance may be observed in some Jumbo frames environments. If this is observed, increasing the application's socket buffer size and/or increasing the /proc/sys/net/ipv4/tcp_*mem entry values may help. For more details, see the specific application documentation in the text file ip-sysctl.txt in your kernel documentation.

Multiple Interfaces on Same Ethernet Broadcast Network

Due to the default ARP behavior on Linux, it is not possible to have one system on two IP networks in the same Ethernet broadcast domain (non-partitioned switch) behave as expected. All Ethernet interfaces will respond to IP traffic for any IP address assigned to the system. This results in unbalanced receive traffic.

If you have multiple interfaces in a server, turn on ARP filtering by entering:

```
echo 1 > /proc/sys/net/ipv4/conf/all/arp_filter
```

(this only works if your kernel's version is higher than 2.4.5), or install the interfaces in separate broadcast domains.

UDP Stress Test Dropped Packet Issue

Under small packets UDP stress test with the 10GbE driver, the Linux system may drop UDP packets due to the fullness of socket buffers. You may want to change the driver's Flow Control variables to the minimum value for controlling packet reception.

Another option is to increase the kernel's default buffer sizes for udp by changing the values in /proc/sys/net/core/rmem_default and rmem_max.

Unplugging Network Cable While ethtool -p is Running

In kernel versions 2.5.50 and later (including 2.6 kernel), unplugging the network cable while ethtool -p is running will cause the system to become unresponsive to keyboard commands, except for control-alt-delete. Restarting the system appears to be the only remedy.

Cisco Catalyst 4948-10GE Switch Running ethtool -g May Cause Switch to Shut Down Ports

82598-based hardware can re-establish link quickly and when connected to some switches, rapid resets within the driver may cause the switch port to become isolated due to "link flap". This is typically indicated by a yellow instead of a green link light. Several operations may cause this problem, such as repeatedly running ethtool commands that cause a reset.

A potential workaround is to use the Cisco IOS command "no errdisable detect cause all" from the Global Configuration prompt which enables the switch to keep the interfaces up, regardless of errors.

MSI-X Issues with Kernels Between 2.6.19 and 2.6.21 (inclusive)

Kernel panics and instability may be observed on any MSI-X hardware if you use irqbalance with kernels between 2.6.19 and 2.6.21. If these types of problems are encountered, you may disable the irqbalance daemon or upgrade to a newer kernel.

Rx Page Allocation Errors

Page allocation failure order: 0 errors may occur under stress with kernels 2.6.25 and above. This is caused by the way the Linux kernel reports this stressed condition.

LRO and iSCSI Incompatibility

LRO is incompatible with iSCSI target or initiator traffic. A panic may occur when iSCSI traffic is received through the ixgbe driver with LRO enabled. To avoid this issue, the driver should be built and installed with the following:

```
# make CFLAGS_EXTRA=-DIXGBE_NO_LRO install
```

DCB: Generic Segmentation Offload Enabled Causes Bandwidth Allocation Issues

In order for DCB to work correctly, GSO (Generic Segmentation Offload aka software TSO) must be disabled using ethtool. By default since the hardware supports TSO (hardware offload of segmentation) GSO will not be running. The GSO state can be queried with ethtool using ethtool -k ethX.

Disable GRO When Routing/Bridging

Due to a known kernel issue, GRO must be turned off when routing/bridging. GRO can be turned off via ethtool.

```
ethtool -K ethX gro off
```

where ethX is the Ethernet interface you're trying to modify.

Lower than Expected Performance on Dual Port and Quad Port 10 Gigabit Ethernet Devices

Some PCIe x8 slots are actually configured as x4 slots. These slots have insufficient bandwidth for full 10GbE line rate with dual port and quad port 10GbE devices. In addition, if you put a PCIe Gen 3-capable adapter into a PCIe Gen 2 slot, you can not get full bandwidth. The driver can detect this situation and will write the following message in the system log: “PCI Express bandwidth available for this card is not sufficient for optimal performance. For optimal performance a x8 PCI Express slot is required.”

If this error occurs, moving your adapter to a true x8 slot will resolve the issue.

ethtool May Incorrectly Display SFP+ Fiber Module as Direct Attached Cable

Due to kernel limitations, port type can only be correctly displayed on kernel 2.6.33 or greater.

Running ethtool -t ethX Command Causes Break between PF and Test Client

When there are active VFs, "ethtool -t" will only run the link test. The driver will also log in syslog that VFs should be shut down to run a full diagnostics test.

Enabling SR-IOV in a 32-bit or 64-bit Microsoft* Windows* Server 2008/R2 Guest OS Using Intel® X540-based 10GbE Controller or Intel® 82599-based 10GbE Controller Under KVM

KVM Hypervisor/VMM supports direct assignment of a PCIe device to a VM. This includes traditional PCIe devices, as well as SR-IOV-capable devices using Intel X540-based and 82599-based controllers.

While direct assignment of a PCIe device or an SR-IOV Virtual Function (VF) to a Linux-based VM running 2.6.32 or later kernel works fine, there is a known issue with Microsoft Windows Server 2008/R2 VM that results in a "yellow bang" error. This problem is within the KVM VMM itself, not the Intel driver, or the SR-IOV logic of the VMM, but rather that KVM emulates an older CPU model for the guests, and this older CPU model does not support MSI-X interrupts, which is a requirement for Intel SR-IOV.

If you wish to use the Intel X540 or 82599-based controllers in SR-IOV mode with KVM and a Microsoft Windows Server 2008/R2 guest try the following workaround. The workaround is to tell KVM to emulate a different model of CPU when using qemu to create the KVM guest:

```
"-cpu qemu64,model=13"
```

Unable to Obtain DHCP Lease on Boot with RedHat

For configurations where the auto-negotiation process takes more than 5 seconds, the boot script may fail with the following message:

```
"ethX: failed. No link present. Check cable?"
```

If this error appears even though the presence of a link can be confirmed using `ethtool ethX`, set `LINKDELAY=15` in /etc/sysconfig/network-scripts/ifcfg-ethX.

NOTE: Link time can take up to 30 seconds. Adjust the LINKDELAY value accordingly.

Alternatively NetworkManager can be used to configure the interfaces, which avoids the set timeout. For configuration instructions of NetworkManager refer to the documentation provided by your distribution.

Loading ixgbe Driver in 3.2.x and Above Kernels Displays Kernel Tainted Message

Due to recent kernel changes, loading an out of tree driver will cause the kernel to be tainted.

Host may Reboot after Removing PF when VF is Active in Guest

If you are using kernel versions earlier than 3.2, do not unload the PF driver with active VFs. Doing this will cause your VFs to stop working until you reload the PF driver and may cause a spontaneous reboot of your system.

Software Bridging Does Not Work with SR-IOV Virtual Functions

SR-IOV Virtual Functions are unable to send or receive traffic between VMs using emulated connections on a Linux Software bridge and connections that use SR-IOV VFs.

Unloading Physical Function (PF) Driver Causes System Reboots when VM is Running and VF is Loaded on the VM

On pre-3.2 Linux kernels unloading the Physical Function (PF) driver causes system reboots when the VM is running and VF is loaded on the VM.

Do not unload the PF driver (ixgbe) while VFs are assigned to guests.

ixgbevf Linux* Driver for the Intel® 10 Gigabit Server Adapters

SR-IOV Overview

Single Root IO Virtualization (SR-IOV) is a PCI SIG specification allowing PCI Express devices to appear as multiple separate physical PCI Express devices. SR-IOV allows efficient sharing of PCI devices among Virtual Machines (VMs). It manages and transports data without the use of a hypervisor by providing independent memory space, interrupts, and DMA streams for each virtual machine.

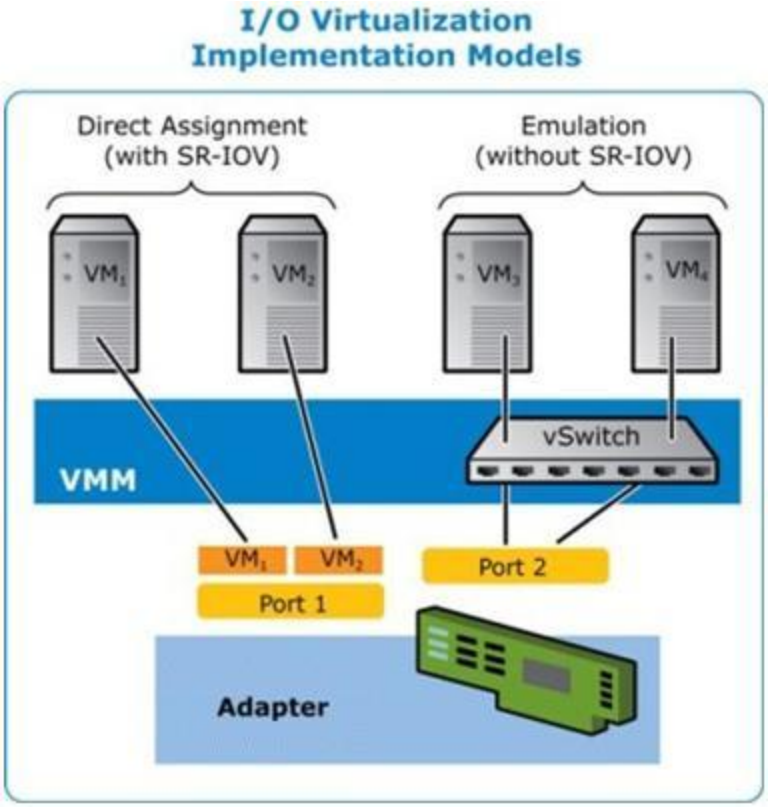

SR-IOV architecture includes two functions:

- Physical Function (PF) is a full featured PCI Express function that can be discovered, managed and configured like any other PCI Express device.
- Virtual Function (VF) is similar to PF but cannot be configured and only has the ability to transfer data in and out. The VF is assigned to a Virtual Machine.

SR-IOV Benefits

SR-IOV has the ability to increase the number of virtual machines supported per physical host, improving I/O device sharing among virtual machines for higher overall performance:

- Provides near native performance due to direct connectivity to each VM through a virtual function
- Preserves VM migration
- Increases VM scalability on a virtualized server
- Provides data protection

SR-IOV Software Requirements

- ixgbe driver - Intel® Linux Base Driver for 82599 and X540-based 10 Gigabit Family of Adapters
- ixgbevf driver - Intel® Linux Driver for 82599 and X540-based 10 Gigabit Family of Adapters
- KVM driver

NOTE: SR-IOV must be enabled in the BIOS.

ixgbevf Driver

SR-IOV is supported by the ixgbevf driver, which should be loaded on both the host and VMs. This driver supports upstream kernel versions 2.6.30 (or higher) x86_64.

The ixgbevf driver supports 82599-based and X540-based virtual function devices that can only be activated on kernels supporting SR-IOV. SR-IOV requires the correct platform and OS support.

The ixgbevf driver requires the ixgbe driver, version 2.0 or later. The ixgbevf driver supports virtual functions generated by the ixgbe driver with a max_vfs value of 1 or greater. For more information on the max_vfs parameter refer to the section on the ixgbe driver.

The guest OS loading the ixgbevf driver must support MSI-X interrupts.

This driver is only supported as a loadable module at this time. Intel is not supplying patches against the kernel source to allow for static linking of the driver. For questions related to hardware requirements, refer to the documentation supplied with your Intel 10GbE adapter. All hardware requirements listed apply to use with Linux.

Instructions on updating ethtool can be found in the section Additional Configurations later in this document.

ixgbevf Linux Base Driver Supported Adapters

The following Intel network adapters are compatible with the ixgbevf Linux driver in this release and can support up to 63 virtual functions per port.

- Intel® Ethernet X520 10GbE Dual Port KX4-KR Mezz
- Intel® Ethernet 10G 2P X540-t Adapter
- Intel® Ethernet 10G 2P X550-t Adapter
- Intel® Ethernet 10G 4P X540/I350 rNDC
- Intel® Ethernet 10G 4P X520/I350 rNDC
- Intel® Ethernet 10G 2P X520-k bNDC
- Intel® Ethernet 10G 2P X520 Adapter
- Intel® Ethernet 10G X520 LOM

For information on how to identify your adapter or for the latest network drivers for Linux, see Customer Support.

SR-IOV Capable Operating Systems

- Citrix XenServer 6.0 with Red Hat Enterprise Linux
- Red Hat* Enterprise Linux* (RHEL) 7.2
- RHEL 7.1
- RHEL 6.7
- Novell* SUSE* Linux Enterprise Server (SLES) 12 SP1
- SLES 11 SP4

Building and Installation

To enable SR-IOV on your system:

1. Ensure both Virtualization and SR-IOV are enabled in the BIOS.
2. Install the Linux operating system. You can verify that the KVM driver is loaded by typing: `lsmod | grep -i kvm`
3. Load the Linux Base Driver using the modprobe command: `modprobe ixgbe option max_vfs=xx,yy`
   xx and *yy* are the number of virtual functions you want to create. You must specify a number for each port with each parameter separated by a comma. For example, *xx* is the number of virtual functions for port 1; and *yy*, for port 2. You can create up to 63 functions per port.
4. Compile and install the ixgbevf driver for SR-IOV. This is loaded against the virtual functions created.

Instructions to configure virtual functions are provided in the Technical Briefs listed below. It should be noted that the information contained in these documents references typical configurations. Consult your operating system vendor for the latest information.

- Using Intel® Ethernet and the PCISIG Single Root I/O Virtualization (SR-IOV) and Sharing Specification on Red Hat Enterprise Linux
- How to Configure Intel® X520 Ethernet Server Adapter Based Virtual Functions on Citrix XenServer 6.0

NOTE: For VLANs, there is a limit of a total of 32 shared VLANs to 1 or more virtual functions.

There are three methods for installing the Linux driver:

- Install from Source Code
- Install from a DKMS RPM
- Install Using KMP RPM

Install from Source Code

To build a binary RPM* package of this driver, run 'rpmbuild -tb <filename.tar.gz>'. Replace <filename.tar.gz> with the specific filename of the driver.

NOTES:

- For the build to work properly it is important that the currently running kernel MATCH the version and configuration of the installed kernel source. If you have just recompiled your kernel, reboot the system.
- RPM functionality has only been tested in Red Hat distributions.

1. Download the base driver tar file to the directory of your choice. For example, use '`/home/username/ixgbevf`' or '`/usr/local/src/ixgbevf`'.
2. Untar/unzip the archive, where <x.x.x> is the version n umber for the driver tar:

   ```
   tar zxf ixgbevf-<x.x.x>.tar.gz
   ```

3. Change to the driver src directory, where `<x.x.x>` is the version number for the driver tar:

   ```
   cd ixgbevf-<x.x.x>/src/
   ```

4. Compile the driver module:

   ```
   make install
   ```

   The binary will be installed as: `/lib/modules/<KERNEL VERSION>/kernel/drivers/net/ixgbevf/ixgbevf.ko`

   The install locations listed above are the default locations. This might differ for various Linux distributions. For more information, see the ldistrib.txt file included in the driver tar.
5. Install the module using the modprobe command for kernel 2.6.x:

   ```
   modprobe ixgbevf <parameter>=<value>
   ```

   For 2.6 based kernels, make sure that the older ixgbevf drivers are removed from the kernel, before loading the new module:

   ```
   rmmod ixgbevf; modprobe ixgbevf
   ```

6. Assign an IP address to and activate the Ethernet interface by entering the following, where <x> is the interface number:

   ```
   ifconfig eth<x> <IP_address> netmask <netmask>
   ```

7. Verify that the interface works. Enter the following, where <IP_address> is the IP address for another machine on the same subnet as the interface that is being tested:

   ```
   ping <IP_address>
   ```

Install from a DKMS RPM

Download the DKMS RPM file. To install or uninstall the RPMs, follow the instructions below.

NOTE: The DKMS (2.0 or higher) framework must be installed before attempting to install a DKMS RPM. DKMS can be acquired at: http://linux.dell.com/dkms/

DKMS RPMs are provided for supported Linux distributions and are named ixgbevf-<driver_version>-<type>.noarch.rpm

<driver_version> and <type> are specific to the driver version you are using:

- <driver_version> is the driver version number; e.g., 6.2.xy.
- <type> is either "sb_dkms" for packages that contain both source and binaries, or "bo_dkms" for packages that contain binary-only modules.

The naming convention for the included DKMS RPMs is:

```
<component name>-<component version>-<type>.noarch.rpm
```

For example, ixgbevf-x.y.z-sb_dkms.noarch.rpm: ixgbevf is the component name; x.y.z is the component version; and sb_dkms is the type.

NOTE: If you have a previous RPM installed, it must be uninstalled before installing the new RPM. To determine whether or not a previous RPM is installed, enter rpm -q ixgbevf.

To install the DKMS RPM, use this command:

```
rpm -i <filename>
```

To uninstall the RPM, use this command:

```
rpm -e <package>
```

Install Using KMP RPM

NOTE: KMP is only supported on SLES11.

The KMP RPMs update existing ixgbevf RPMs currently installed on the system. These updates are provided by SuSE in the SLES release. If an RPM does not currently exist on the system, the KMP will not install.

The RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

intel-<component name>-<component version>.<arch type>.rpm

For example, intel-ixgbevf-1.3.8.6-1.x86_64.rpm: ixgbevf is the component name; 1.3.8.6-1 is the component version; and x86_64 is the architecture type.

KMP RPMs are provided for supported Linux distributions. The naming convention for the included KMP RPMs is:

intel-<component name>-kmp-<kernel type>-<component version>_<kernel version>.<arch type>.rpm

For example, intel-ixgbevf-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm: ixgbevf is the component name; default is the kernel type; 1.3.8.6 is the component version; 2.6.27.19_5-1 is the kernel version; and x86_64 is the architecture type.

To install the KMP RPM, type the following two commands:

```
rpm -i <rpm filename>
rpm -i <kmp rpm filename>
```

For example, to install the ixgbevf KMP RPM package, type the following:

```
rpm -i intel-ixgbevf-1.3.8.6-1.x86_64.rpm
rpm -i intel-ixgbevf-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm
```

Command Line Parameters

If the driver is built as a module, the following optional parameters are used by entering them on the command line with the modprobe command using this syntax:

```
modprobe ixgbevf [<option>=<VAL1>,<VAL2>,...]
```

For example:

```
modprobe ixgbevf InterruptThrottleRate=16000,16000
```

The default value for each parameter is generally the recommended setting, unless otherwise noted.

The following table contains parameters and possible values for modprobe commands:

Parameter Name	Valid Range/Settings	Default	Description
InterruptThrottleRate	0, 1, 956 - 488,281 (0=off, 1=dynamic)	8000	The driver can limit the number of interrupts per second that the adapter will generate for incoming packets. It does this by writing a value to the adapter that is based on the maximum number of interrupts that the adapter will generate per second. Setting InterruptThrottleRate to a value greater or equal to 100 will program the adapter to send out a maximum of that many interrupts per second, even if more packets have come in. This reduces interrupt load on the system and can lower CPU utilization under heavy load, but will increase latency as packets are not processed as quickly. The default behavior of the driver previously assumed a static InterruptThrottleRate value of 8000, providing a good fallback value for all traffic types, but lacking in small packet performance and latency. The hardware can handle many more small packets per second however, and for this reason an adaptive interrupt moderation algorithm was implemented. The driver has one adaptive mode (setting 1) in which it dynamically adjusts the InterruptThrottleRate value based on the traffic that it receives. After determining the type of incoming traffic in the last time frame, it will adjust the InterruptThrottleRate to an appropriate value for that traffic. The algorithm classifies the incoming traffic every interval into classes. Once the class is determined, the InterruptThrottleRate value is adjusted to suit that traffic type the best. There are three classes defined: "Bulk traffic", for large amounts of packets of normal size; "Low latency", for small amounts of traffic and/or a significant percentage of small packets; and "Lowest latency", for almost completely small packets or minimal traffic. In dynamic conservative mode, the InterruptThrottleRate value is set to 4000 for traffic that falls in class "Bulk traffic". If traffic falls in the "Low latency" or "Lowest latency" class, the InterruptThrottleRate is increased stepwise to 20000. This default mode is suitable for most applications. For situations where low latency is vital such as cluster or grid computing, the algorithm can reduce latency even more when InterruptThrottleRate is set to mode 1. In this mode, the InterruptThrottleRate will be increased stepwise to 70000 for traffic in class "Lowest latency".

Parameter Name	Valid Range/Settings	Default	Description
			Setting InterruptThrottleRate to 0 turns off any interrupt moderation and may improve small packet latency, but is generally not suitable for bulk throughput traffic. **NOTES:** • Dynamic interrupt throttling is only applicable to adapters operating in MSI or Legacy interrupt mode, using a single receive queue. • When ixgbevf is loaded with default settings and multiple adapters are in use simultaneously, the CPU utilization may increase non-linearly. In order to limit the CPU utilization without impacting the overall throughput, it is recommended to load the driver as follows `modprobe ixgbevf InterruptThrottleRate=3000,3000,3000` This sets the InterruptThrottleRate to 3000 interrupts/sec for the first, second, and third instances of the driver. The range of 2000 to 3000 interrupts per second works on a majority of systems and is a good starting point, but the optimal value will be platform-specific. If CPU utilization is not a concern, use default driver settings.

NOTES:

- For more information about the InterruptThrottleRate parameter, see the application note at http://www.intel.com/design/network/applnots/ap450.htm.
- A descriptor describes a data buffer and attributes related to the data buffer. This information is accessed by the hardware.

Additional Configurations

Configuring the Driver on Different Distributions

Configuring a network driver to load properly when the system is started is distribution dependent. Typically, the configuration process involves adding an alias line to /etc/modules.conf or /etc/modprobe.conf as well as editing other system startup scripts and/or configuration files. Many Linux distributions ship with tools to make these changes for you. To learn the proper way to configure a network device for your system, refer to your distribution documentation. If during this process you are asked for the driver or module name, the name for the Linux Base Driver for the Intel® 10 Gigabit PCI Express Family of Adapters is ixgbevf.

Viewing Link Messages

Link messages will not be displayed to the console if the distribution is restricting system messages. In order to see network driver link messages on your console, set dmesg to eight by entering the following:

```
dmesg -n 8
```

NOTE: This setting is not saved across reboots.

ethtool

The driver uses the ethtool interface for driver configuration and diagnostics, as well as displaying statistical information. The latest ethtool version is required for this functionality.

The latest release of ethtool can be found at: http://sourceforge.net/projects/gkernel.

MACVLAN

ixgbevf supports MACVLAN on those kernels that have the feature included. Kernel support for MACVLAN can be tested by checking if the MACVLAN driver is loaded. The user can run '`lsmod | grep macvlan`' to see if the MACVLAN driver is loaded or run '`modprobe macvlan`' to try to load the MACVLAN driver.

It may be necessary to update to a recent release of the iproute2 package to get support of MACVLAN via the 'ip' command.

NAPI

NAPI (Rx polling mode) is supported in the ixgbevf driver and is always enabled. For more information on NAPI, go to: ftp://robur.slu.se/pub/Linux/net-development/NAPI/usenix-paper.tgz.

Known Issues

NOTE: After installing the driver, if your Intel network connection is not working, verify that you have installed the correct driver.

Driver Compilation

When trying to compile the driver by running make install, the following error may occur: "Linux kernel source not configured - missing version.h"

To solve this issue, create the version.h file by going to the Linux source tree and entering:

```
make include/linux/version.h
```

Multiple Interfaces on Same Ethernet Broadcast Network

Due to the default ARP behavior on Linux, it is not possible to have one system on two IP networks in the same Ethernet broadcast domain (non-partitioned switch) behave as expected. All Ethernet interfaces will respond to IP traffic for any IP address assigned to the system. This results in unbalanced receive traffic.

If you have multiple interfaces in a server, turn on ARP filtering by entering:

```
echo 1 > /proc/sys/net/ipv4/conf/all/arp_filter
```

(this only works if your kernel's version is higher than 2.4.5), or install the interfaces in separate broadcast domains.

NOTE: This setting is not saved across reboots. The configuration change can be made permanent by adding the line:

`net.ipv4.conf.all.arp_filter= 1` to the file /etc/sysctl.conf

or

install the interfaces in separate broadcast domains (either in different switches or in a switch partitioned to VLANs).

MSI-X Issues with Kernels Between 2.6.19 and 2.6.21 (inclusive)

Kernel panics and instability may be observed on any MSI-X hardware if you use irqbalance with kernels between 2.6.19 and 2.6.21. If these types of problems are encountered, you may disable the irqbalance daemon or upgrade to a newer kernel.

Rx Page Allocation Errors

Page allocation failure order:0 errors may occur under stress with kernels 2.6.25 and above. This is caused by the way the Linux kernel reports this stressed condition.

Host may Reboot after Removing PF when VF is Active in Guest

If you are using kernel versions earlier than 3.2, do not unload the PF driver with active VFs. Doing this will cause your VFs to stop working until you reload the PF driver and may cause a spontaneous reboot of your system.

i40e Linux* Driver for the Intel X710 Ethernet Controller Family

i40e Overview

The i40e Linux* Base Driver for the X710/XL710 Ethernet Controller Family of Adapters supports the 2.6.32 and newer kernels, and includes support for Linux supported x86_64 systems. For questions related to hardware requirements, refer to System Requirements. All hardware requirements listed apply to use with Linux.

The following features are now available in supported kernels:

- VXLAN encapsulation
- Native VLANs
- Channel Bonding (teaming)
- SNMP
- Generic Receive Offload
- Data Center Bridging

Adapter teaming is now implemented using the native Linux Channel bonding module. This is included in supported Linux kernels. Channel Bonding documentation can be found in the Linux kernel source: /Documentation/networking/bonding.txt

Use ethtool, lspci, or iproute2's ip command to obtain driver information. Instructions on updating ethtool can be found in the Additional Configurations section later in this page.

i40e Linux Base Driver Supported Devices

The following Intel network adapters are compatible with this driver:

- Intel® Ethernet 10G 4P X710-k bNDC
- Intel® Ethernet 10G 2P X710-k bNDC
- Intel® Ethernet 10G X710-k bNDC
- Intel® Converged Network Adapter X710
- Intel® Ethernet 10G 4P X710/l350 rNDC
- Intel® Ethernet 10G 4P X710 SFP+ rNDC
- Intel® Ethernet 10G X710 rNDC
- Intel® Ethernet 40G 2P XL710 QSFP+ rNDC
- Intel® Ethernet Converged Network Adapter XL710-Q2

For information on how to identify your adapter or for the latest network drivers for Linux, see Customer Support.

SFP+ Devices with Pluggable Optics

NOTE: For SFP+ fiber adapters, using "`ifconfig down`" turns off the laser. "`ifconfig up`" turns the laser on.

For information on using SFP+ devices with pluggable optics, click here.

Building and Installation

There are four methods for installing the Linux driver:

- Install from Source Code
- Install Using KMP RPM
- Install Using KMOD RPM

Install from Source Code

To build a binary RPM* package of this driver, run 'rpmbuild -tb <filename.tar.gz>'. Replace <filename.tar.gz> with the specific filename of the driver.

NOTES:

- For the build to work properly it is important that the currently running kernel MATCH the version and configuration of the installed kernel source. If you have just recompiled your kernel, reboot the system.
- RPM functionality has only been tested in Red Hat distributions.

1. Download the base driver tar file to the directory of your choice. For example, use '`/home/username/i40e`' or '`/usr/local/src/i40e`'.
2. Untar/unzip the archive, where <x.x.x> is the version number for the driver tar:

   ```
   tar zxf i40e-<x.x.x>.tar.gz
   ```

3. Change to the driver src directory, where `<x.x.x>` is the version number for the driver tar:

   ```
   cd i40e-<x.x.x>/src/
   ```

4. Compile the driver module:

   ```
   make install
   ```

   The binary will be installed as: `/lib/modules/<KERNEL VERSION>/kernel/drivers/net/i40e/i40e.ko`

   The install locations listed above are the default locations. This might differ for various Linux distributions. For more information, see the ldistrib.txt file included in the driver tar.
5. Install the module using the modprobe command:

   ```
   modprobe i40e <parameter>=<value>
   ```

   Make sure that the older i40e drivers are removed from the kernel, before loading the new module:

   ```
   rmmod i40e; modprobe i40e
   ```

6. Assign an IP address to and activate the Ethernet interface by entering the following, where <ethx> is the interface name:

   ```
   ifconfig <ethx> <IP_address> netmask <netmask> up
   ```

7. Verify that the interface works. Enter the following, where <IP_address> is the IP address for another machine on the same subnet as the interface that is being tested:

   ```
   ping <IP_address>
   ```

Install Using KMP RPM

NOTE: KMP is only supported on RHEL 6.0 and newer, and SLES11, and newer.

The KMP RPMs update existing i40e RPMs currently installed on the system. These updates are provided by SuSE in the SLES release. If an RPM does not currently exist on the system, the KMP will not install.

The RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

intel-<component name>-<component version>.<arch type>.rpm

For example, intel-i40e-1.3.8.6-1.x86_64.rpm:i40e is the component name; 1.3.8.6-1 is the component version; and x86_64 is the architecture type.

KMP RPMs are provided for supported Linux distributions. The naming convention for the included KMP RPMs is:

intel-<component name>-kmp-<kernel type>-<component version>_<kernel version>.<arch type>.rpm

For example, intel-i40e-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm:i40e is the component name; default is the kernel type; 1.3.8.6 is the component version; 2.6.27.19_5-1 is the kernel version; and x86_64 is the architecture type.

To install the KMP RPM, type the following two commands:

```
rpm -i <rpm filename>
rpm -i <kmp rpm filename>
```

For example, to install the i40e KMP RPM package, type the following:

```
rpm -i intel-i40e-1.3.8.6-1.x86_64.rpm
rpm -i intel-i40e-kmp-default-1.3.8.6_2.6.27.19_5-1.x86_64.rpm
```

Install Using KMOD RPM

The KMOD RPMs are provided for supported Linux distributions. The naming convention for the included RPMs is:

```
kmod-<driver name>-<version>-1.<arch type>.rpm
```

For example, kmod-i40e-2.3.4-1.x86_64.rpm:

- i40e is the driver name
- 2.3.4 is the version
- x86_64 is the architecture type

To install the KMOD RPM, go to the directory of the RPM and type the following command:

```
rpm -i <rpm filename>
```

For example, to install the i40e KMOD RPM package from RHEL 6.4, type the following:

```
rpm -i kmod-i40e-2.3.4-1.x86_64.rpm
```

Command Line Parameters

In general, ethtool and other OS specific commands are used to configure user changeable parameters after the driver is loaded. The i40e driver only supports the max_vfs kernel parameter on older kernels that do not have the standard sysfs interface. The only other module parameter is the debug parameter that can control the default logging verbosity of the driver.

If the driver is built as a module, the following optional parameters are used by entering them on the command line with the modprobe command using this syntax:

```
modprobe i40e [<option>=<VAL1>]
```

For example:

```
modprobe i40e max_vfs=7
```

The default value for each parameter is generally the recommended setting, unless otherwise noted.

The following table contains parameters and possible values for modprobe commands:

Parameter Name	Valid Range/Settings	Default	Description

max_vfs	1 - 63	0	This parameter adds support for SR-IOV. It causes the driver to spawn up to max_vfs worth of virtual function.
			
			NOTES:
			• On kernel version 3.8.x and above VFs are created by writing an appropriate value to the sriov_numvfs parameter via the sysfs interface.
			• When either SR-IOV mode is enabled, hardware VLAN filtering and VLAN tag stripping/insertion will remain enabled. Please remove the old VLAN filter before the new VLAN filter is added. For example:
			
			`ip link set eth0 vf 0 vlan 100`
			`// set vlan 100 for VF 0`
			`ip link set eth0 vf 0 vlan 0`
			`// Delete vlan 100`
			`ip link set eth0 vf 0 vlan 200`
			`// set a new vlan 200 for VF 0`

Parameter Name	Valid Range/Settings	Default	Description
			With kernel 3.6, the driver supports the simultaneous usage of max_vfs and DCB features, subject to the constraints described below. Prior to kernel 3.6, the driver did not support the simultaneous operation of max_vfs > 0 and the DCB features (multiple traffic classes utilizing Priority Flow Control and Extended Transmission Selection). When DCB is enabled, network traffic is transmitted and received through multiple traffic classes (packet buffers in the NIC). The traffic is associated with a specific class based on priority, which has a value of 0 through 7 used in the VLAN tag. When SR-IOV is not enabled, each traffic class is associated with a set of RX/TX descriptor queue pairs. The number of queue pairs for a given traffic class depends on the hardware configuration. When SR-IOV is enabled, the descriptor queue pairs are grouped into pools. The Physical Function (PF) and each Virtual Function (VF) is allocated a pool of RX/TX descriptor queue pairs. When multiple traffic classes are configured (for example, DCB is enabled), each pool contains a queue pair from each traffic class. When a single traffic class is configured in the hardware, the pools contain multiple queue pairs from the single traffic class. The number of VFs that can be allocated depends on the number of traffic classes that can be enabled. The configurable number of traffic classes for each enabled VF is as follows: 0 - 15 VFs = Up to 8 traffic classes, depending on device support 16 - 31 VFs = Up to 4 traffic classes 32 - 63 = 1 traffic class When VFs are configured, the PF is allocated one pool as well. The PF supports the DCB features with the constraint that each traffic class will only use a single queue pair. When zero VFs are configured, the PF can support multiple queue pairs per traffic class.
Intel® Ethernet Flow Director			**NOTE:** Flow director parameters are only supported on kernel versions 2.6.30 or later. Flow control in 1 Gbps mode on these devices can lead to transmit hangs. This supports advanced filters that direct receive packets by their flows to different queues and enables tight control on routing a flow in the platform. It matches flows and CPU cores for flow affinity and supports multiple parameters for flexible flow classification and load balancing. The flow director is enabled only if the kernel is multiple TX queue capable. An included script (set_irq_affinity.sh) automates setting the IRQ to CPU affinity. To verify that the driver is using Flow Director, look at the counter in ethtool: fdir_miss and fdir_match. Other ethtool Commands: To enable/disable Flow Director `ethtool -K ethX ntuple <on|off>` When disabling ntuple filters all the user programed filters get flushed from the driver cache and HW. The user must re-add the filters if needed when ntuple is re-enabled. To add a filter that directs packet to queue 2, use the -U or -N switch. e.g.,

Parameter Name	Valid Range/Settings	Default	Description

			`# ethtool -N ethX flow-type tcp4 src-ip 192.168.10.1 dst-ip \ 192.168.10.2 src-port 2000 dst-port 2001 action 2 [loc 1]`
			
			To see the list of filters currently present
			
			`# ethtool <-u|-n> ethX`
			
			NOTE: i40e linux driver does not support configuration of the mask field. It only accepts rules that completely qualify a certain flow type.
			
			ATR (application Targeted Routing) Perfect Filter:
			
			ATR is enabled by default when the kernel is in multiple TX queue mode. An ATR flow director filter rule is added when a TCP-IP flow starts and is deleted when the flow ends. When a TCP-IP Flow Director rule is added from ethtool (Sideband filter), ATR is turned off by the driver. In order to re-enable ATR the user can disable side-band by ethtool -K option. If side-band is re-enabled after this, ATR will remain enabled until a TCP-IP flow is added.
			
			Sideband Perfect filters
			
			Sideband Perfect filters is an interface to load the filter table that funnels all flows into queue_0 unless an alternative queue is specified using "action". In that case, any flow that matches the filter criteria will be directed to the appropriate queue.
			
			Rules may be deleted from the table itself. This is done via "ethtool -U ethX delete N" where N is the rule number to be deleted.
			
			If the queue is defined as -1, the filter will drop matching packets. To account for Sideband filter matches, there are fdir_sb_match stats in ethtool.
			
			In addition, rx-N.rx_packets shows the number of packets processed by the Nth queue.
			
			NOTES:
			• Receive Packet Steering (RPS) and Receive Flow Steering (RFS) are not compatible with Flow Director. If Flow Director is enabled, these will be disabled.
			• Once a rule is defined, to delete a rule you must supply the same fields and masks (if masks are specified).
			
			Known Issues
			
			Flow Director Sideband Logic adds duplicate filter in the SW filter list:
			
			The Flow Director Sideband Logic adds a duplicate filter in the SW filter list if the location is not specified or is specified but differs from the previous rule location but has the same filter criteria. The second of the two filters that appear is the valid one in HW and it decides the filter action.

Additional Configurations

Configuring the Driver on Different Distributions

Configuring a network driver to load properly when the system is started is distribution dependent. Typically, the configuration process involves adding an alias line to /etc/modules.conf or /etc/modprobe.conf as well as editing other system startup scripts and/or configuration files. Many Linux distributions ship with tools to make these changes for you. To learn the proper way to configure a network device for your system, refer to your distribution documentation. If during this process you are asked for the driver or module name, the name for the Linux Base Driver for the Intel® 10 Gigabit PCI Express Family of Adapters is i40e.

Viewing Link Messages

Link messages will not be displayed to the console if the distribution is restricting system messages. In order to see network driver link messages on your console, set dmesg to eight by entering the following:

```
dmesg -n 8
```

NOTE: This setting is not saved across reboots.

Jumbo Frames

Jumbo Frames support is enabled by changing the MTU to a value larger than the default of 1500 bytes. The maximum value for the MTU is 9710. Use the ifconfig command to increase the MTU size. For example, enter the following where <x> is the interface number:

```
ifconfig ethx mtu 9000 up
```

This setting is not saved across reboots. The setting change can be made permanent by adding `MTU = 9000` to the file `/etc/sysconfig/network-scripts/ifcfg-eth<x>` for RHEL or to the file `/etc/sysconfig/network/<config_file>` for SLES.

The maximum MTU setting for Jumbo Frames is 9710. This value coincides with the maximum Jumbo Frames size of 9728. This driver will attempt to use multiple page sized buffers to receive each jumbo packet. This should help to avoid buffer starvation issues when allocating receive packets.

ethtool

The driver uses the ethtool interface for driver configuration and diagnostics, as well as displaying statistical information. The latest ethtool version is required for this functionality.

The latest release of ethtool can be found at: http://sourceforge.net/projects/gkernel.

NAPI

NAPI (Rx polling mode) is supported in the i40e driver.

See ftp://robur.slu.se/pub/Linux/net-development/NAPI/usenix-paper.tgz for more information on NAPI.

Flow Control

Flow control is disabled by default. To enable it, use ethtool:

```
ethtool -A eth? autoneg off rx on tx on
```

NOTE: You must have a flow control capable link partner.

MAC and VLAN Anti-spoofing Feature

When a malicious driver attempts to send a spoofed packet, it is dropped by the hardware and not transmitted. An interrupt is sent to the PF driver notifying it of the spoof attempt. When a spoofed packet is detected the PF driver will send the following message to the system log (displayed by the "dmesg" command):

```
i40e ethx: i40e_spoof_check: n spoofed packets detected
```

Where x=the PF interface# and n=the VF that attempted to do the spoofing.

NOTE: This feature can be disabled for a specific Virtual Function (VF).

Support for UDP RSS

This feature adds an ON/OFF switch for hashing over certain flow types. The default setting is disabled. NOTE: Fragmented packets may arrive out of order when RSS UDP support is configured.

Supported ethtool Commands and Options

`-n --show-nfc`
Retrieves the receive network flow classification configurations.

`rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6`
Retrieves the hash options for the specified network traffic type.

`-N --config-nfc`
Configures the receive network flow classification.

`rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6 m|v|t|s|d|f|n|r...`
Configures the hash options for the specified network traffic type.

udp4 UDP over IPv4
udp6 UDP over IPv6

f Hash on bytes 0 and 1 of the Layer 4 header of the rx packet.
n Hash on bytes 2 and 3 of the Layer 4 header of the rx packet.

VXLAN Overlay HW Offloading

The i40e Linux driver features VXLAN Overlay HW Offloading support. The following two commands are used to view and configure VXLAN on a VXLAN-overlay offload enabled device.

This command displays the offloads and their current state:

```
# ethtool -k ethX
```

This command enables/disables VXLAN support in the driver.

```
# ethtool -K ethX tx-udp_tnl-segmentation [off|on]
```

For more information on configuring your network for VXLAN overlay support, refer to the Intel Technical Brief, "Creating Overlay Networks Using Intel Ethernet Converged Network Adapters" (Intel Networking Division, August 2013): http://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/ overlay-networks-using-converged-network-adapters-brief.pdf

NPar

On X710/XL710 based adapters that support it, you can set up multiple functions on each physical port. You configure these functions through the System Setup/BIOS.

Minimum TX Bandwidth is the guaranteed minimum data transmission bandwidth, as a percentage of the full physical port link speed, that the partition will receive. The bandwidth the partition is awarded will never fall below the level you specify here.

The range for the minimum bandwidth values is:

1 to ((100 minus # of partitions on the physical port) plus 1)

For example, if a physical port has 4 partitions, the range would be

1 to ((100 - 4) + 1 = 97)

The Maximum Bandwidth percentage represents the maximum transmit bandwidth allocated to the partition as a percentage of the full physical port link speed. The accepted range of values is 1-100. The value can be used as a limiter, should you chose that any one particular function not be able to consume 100% of a port's bandwidth (should it be available). The sum of all the values for Maximum Bandwidth is not restricted, because no more than 100% of a port's bandwidth can ever be used.

Once the initial configuration is complete, you can set different bandwidth allocations on each function as follows:

1. Make a new directory named /config
2. Edit etc/fstab to include:

   configfs /config configfs defaults
3. Load (or reload) the i40e driver
4. Mount /config
5. Make a new directory under config for each partition upon which you wish to configure the bandwidth.

Three files will appear under the config/partition directory:

- max_bw

- min_bw

- commit

Read from max_bw to get display the current maximum bandwidth setting.

Write to max_bw to set the maximum bandwidth for this function.

Read from min_bw to display the current minimum bandwidth setting.

Write to min_bw to set the minimum bandwidth for this function.

Write a '1' to commit to save your changes.

NOTES:

- commit is write only. Attempting to read it will result in an error.
- Writing to commit is only supported on the first function of a given port. Writing to a subsequent function will result in an error.
- Oversubscribing the minimum bandwidth is not supported. The underlying device's NVM will set the minimum bandwidth to supported values in an indeterminate manner. Remove all of the directories under config and reload them to see what the actual values are.
- To unload the driver you must first remove the directories created in step 5, above.

Example of Setting the minimum and maximum bandwidth (assume there are four function on the port eth6-eth9, and that eth6 is the first function on the port):

```
# mkdir /config/eth6

# mkdir /config/eth7

# mkdir /config/eth8

# mkdir /config/eth9

# echo 50 > /config/eth6/min_bw

# echo 100 > /config/eth6/max_bw

# echo 20 > /config/eth7/min_bw

# echo 100 > /config/eth7/max_bw

# echo 20 > /config/eth8/min_bw

# echo 100 > /config/eth8/max_bw

# echo 10 > /config/eth9/min_bw

# echo 25 > /config/eth9/max_bw
```

```
# echo 1 > /config/eth6/commit
```

Performance Optimization

The driver defaults are meant to fit a wide variety of workloads. If further optimization is required, we recommend experimenting with the following settings.

Pin the adapter's IRQs to specific cores by disabling the irqbalance service and running the included set_irq_affinity script.

The following settings will distribute the IRQs across all the cores evenly:

```
# scripts/set_irq_affinity -x all <interface1> , [ <interface2>,… ]
```

The following settings will distribute the IRQs across all the cores that are local to the adapter (same NUMA node):

```
# scripts/set_irq_affinity -x local <interface1> ,[ <interface2>,… ]
```

Please see the script's help text for further options.

For very CPU intensive workloads, we recommend pinning the IRQs to all cores.

For IP Forwarding: Disable Adaptive ITR and lower rx and tx interrupts per queue using ethtool. Setting rx-usecs and tx-usecs to 125 will limit interrupts to about 8000 interrupts per second per queue.

```
# ethtool <interface> adaptive-rx off adaptive-tx off rx-usecs 125 tx-usecs 125
```

For lower CPU utilization: Disable Adaptive ITR and lower rx and tx interrupts per queue using ethtool. Setting rx-usecs and tx-usecs to 250 will limit interrupts to about 4000 interrupts per second per queue.

```
# ethtool <interface> adaptive-rx off adaptive-tx off rx-usecs 250 tx-usecs 250
```

For lower latency: Disable Adaptive ITR and ITR by setting rx and tx to 0 using ethtool.

```
# ethtool <interface> adaptive-rx off adaptive-tx off rx-usecs 0 tx-usecs 0
```

Known Issues

Driver Compilation

When trying to compile the driver by running make install, the following error may occur: "Linux kernel source not configured - missing version.h"

To solve this issue, create the version.h file by going to the Linux source tree and entering:

```
make include/linux/version.h
```

Performance Degradation with Jumbo Frames

Degradation in throughput performance may be observed in some Jumbo frames environments. If this is observed, increasing the application's socket buffer size and/or increasing the /proc/sys/net/ipv4/tcp_*mem entry values may help. For more details, see the specific application documentation in the text file ip-sysctl.txt in your kernel documentation.

Multiple Interfaces on Same Ethernet Broadcast Network

Due to the default ARP behavior on Linux, it is not possible to have one system on two IP networks in the same Ethernet broadcast domain (non-partitioned switch) behave as expected. All Ethernet interfaces will respond to IP traffic for any IP address assigned to the system. This results in unbalanced receive traffic.

If you have multiple interfaces in a server, turn on ARP filtering by entering:

```
 echo 1 > /proc/sys/net/ipv4/conf/all/arp_filter
```

(this only works if your kernel's version is higher than 2.4.5), or install the interfaces in separate broadcast domains.

UDP Stress Test Dropped Packet Issue

Under small packets UDP stress test with the i40e driver, the Linux system may drop UDP packets due to the fullness of socket buffers. You may want to change the driver's Flow Control variables to the minimum value for controlling packet reception.

Another option is to increase the kernel's default buffer sizes for udp by changing the values in /proc/sys/net/-core/rmem_default and rmem_max.

Unplugging Network Cable While ethtool -p is Running

In kernel versions 2.6.32 and later, unplugging the network cable while ethtool -p is running will cause the system to become unresponsive to keyboard commands, except for control-alt-delete. Restarting the system appears to be the only remedy.

Rx Page Allocation Errors

Page allocation failure order: 0 errors may occur under stress with kernels 2.6.25 and above. This is caused by the way the Linux kernel reports this stressed condition.

Disable GRO When Routing/Bridging

Due to a known kernel issue, GRO must be turned off when routing/bridging. GRO can be turned off via ethtool.

```
ethtool -K ethX gro off
```

where ethX is the Ethernet interface you're trying to modify.

Lower than Expected Performance

Some PCIe x8 slots are actually configured as x4 slots. These slots have insufficient bandwidth for full 10GbE line rate with dual port and quad port 10GbE devices. In addition, if you put a PCIe Gen 3-capable adapter into a PCIe Gen 2 slot, you can not get full bandwidth. The driver can detect this situation and will write the following message in the system log: “PCI Express bandwidth available for this card is not sufficient for optimal performance. For optimal performance a x8 PCI Express slot is required.”

If this error occurs, moving your adapter to a true x8 slot will resolve the issue.

ethtool May Incorrectly Display SFP+ Fiber Module as Direct Attached Cable

Due to kernel limitations, port type can only be correctly displayed on kernel 2.6.33 or greater.

Running ethtool -t ethX Command Causes Break between PF and Test Client

When there are active VFs, "ethtool -t" will only run the link test. The driver will also log in syslog that VFs should be shut down to run a full diagnostics test.

Enabling SR-IOV in a 32-bit or 64-bit Microsoft* Windows* Server 2008/R2 Guest OS Under Linux KVM

KVM Hypervisor/VMM supports direct assignment of a PCIe device to a VM. This includes traditional PCIe devices, as well as SR-IOV-capable devices using Intel X540-based and 82599-based controllers.

While direct assignment of a PCIe device or an SR-IOV Virtual Function (VF) to a Linux-based VM running 2.6.32 or later kernel works fine, there is a known issue with Microsoft Windows Server 2008/R2 VM that results in a "yellow bang" error. This problem is within the KVM VMM itself, not the Intel driver, or the SR-IOV logic of the VMM, but rather that KVM emulates an older CPU model for the guests, and this older CPU model does not support MSI-X interrupts, which is a requirement for Intel SR-IOV.

If you wish to use the Intel X540 or 82599-based controllers in SR-IOV mode with KVM and a Microsoft Windows Server 2008/R2 guest try the following workaround. The workaround is to tell KVM to emulate a different model of CPU when using qemu to create the KVM guest:

```
"-cpu qemu64,model=13"
```

Unable to Obtain DHCP Lease on Boot with RedHat

For configurations where the auto-negotiation process takes more than 5 seconds, the boot script may fail with the following message:

```
"ethX: failed. No link present. Check cable?"
```

If this error appears even though the presence of a link can be confirmed using `ethtool ethX`, set `LINKDELAY=15` in /etc/sysconfig/network-scripts/ifcfg-ethX.

NOTE: Link time can take up to 30 seconds. Adjust the LINKDELAY value accordingly.

Alternatively NetworkManager can be used to configure the interfaces, which avoids the set timeout. For configuration instructions of NetworkManager refer to the documentation provided by your distribution.

Loading i40e Driver in 3.2.x and Above Kernels Displays Kernel Tainted Message

Due to recent kernel changes, loading an out of tree driver will cause the kernel to be tainted.

Software Bridging Does Not Work with SR-IOV Virtual Functions

SR-IOV Virtual Functions are unable to send or receive traffic between VMs using emulated connections on a Linux Software bridge and connections that use SR-IOV VFs.

Troubleshooting

Common Problems and Solutions

Problem	Solution
Your server cannot find the adapter.	• Make sure the adapter is seated firmly in the slot. • Try a different PCI Express* slot. • Use the drivers that came with your adapter or download the latest ones from Customer Support. • Check to see if your motherboard has the latest BIOS. • Try rebooting the server. • Try a different Intel adapter. • Make sure the adapter slot is compatible with the type of adapter being used.
Diagnostics pass but the connection fails.	• Check the responding link partner. • Make sure the cable is securely attached, is the proper type and does not exceed the recommended lengths. • Try another cable. • Try running the Sender-Responder diagnostic test. • Make sure the duplex mode and speed setting on the adapter matches the setting on the switch.
Another adapter stopped working after you installed a second adapter.	• Make sure the cable connections are correct. • Make sure your BIOS is current. • Check for interrupt conflicts and sharing problems. Make sure the other adapter supports shared interrupts. Also, make sure your operating system supports shared interrupts. • Unload all PCI Express device drivers, then reload all devices. • Try reseating all adapters.
The adapter stopped working without apparent cause.	• The network driver files may be damaged or deleted. Reinstall the drivers. • Try reseating the adapter in its slot or a different slot, if necessary. • Try rebooting the server. • Try another cable. • Try a different network adapter. • Run the adapter and network tests described under Testing the Adapter.
The link indicator light is off.	• Run the adapter and network tests described in Testing the Adapter. • Make sure you have loaded the adapter driver. • Check all connections at the adapter and the switch. • Try another port on the switch. • Make sure the cable is securely attached. Also make sure that it is the proper type and does not exceed the recommended lengths. • Make sure the link partner is configured to auto-negotiate (or forced to match adapter). • Verify that your switch is compatible with your network adapter port.
The link light is on, but communications are not	• Make sure the proper (and latest) driver is loaded.

Problem	Solution
properly established.	• Both the adapter and its link partner must be set to either auto-negotiate or manually set to the same speed and duplex settings. **NOTE:** The adapter's link indicator light may be on even if communications between the adapter and its link partner have not been properly established. Technically, the link indicator light represents the presence of a carrier signal but not necessarily the ability to properly communicate with a link partner. This is expected behavior and is consistent with IEEE's specification for physical layer operation.
The diagnostic utility reports the adapter is "Not enabled by BIOS."	• The BIOS isn't configuring the adapter correctly. • Try another PCI Express slot.
The server hangs when the drivers are loaded.	• Change the BIOS interrupt settings.
Adapter unable to connect to switch at correct speed (copper wiring only).	*This is applicable only to copper-based connections.* • Try another cable. • Make sure the cable is Category 5 or 5e. • Make sure the link partner is set to auto-negotiate. • Verify that your switch is compatible with your network adapter port and that you are running the latest operating system revision for your switch.
The device does not connect at the expected speed.	When Gigabit Master/Slave mode is forced to "master" mode on both the Intel adapter and its link partner, the link speed obtained by the Intel adapter may be lower than expected.
After upgrading operating systems, Intel PROSet is no longer available.	If you are upgrading Windows operating systems and you have Intel PROSet software installed, it will be deleted in the process. You will need to reinstall Intel PROSet.
Hardware tests fail in Intel PROSet, DIAGS, Dell OpenManage Server Administrator, or ethtool.	A component on the adapter is not functioning correctly. Try another adapter. If the test fails with the replacement adapter, there might be problem with the system board, or another device on the bus might not be functioning correctly. Contact Customer Support for more information.
The Link test fails in Intel PROSet, DIAGS, Dell OpenManage Server Administrator, or ethtool.	• Make sure the network cable is plugged into the adapter and link partner. • Try connecting to a different port on the link partner. • Try attaching a different cable. • For gigabit copper connections, make sure the cable is a 4-pair Category 5 or better. • For 10 gigabit copper connections, make sure the cable is a 4-pair Category 6 or better. • Make sure the cable length does not exceed specifications. • If the cable is secure, try attaching the cable to a different port on the link partner or try attaching a different cable.
The Connection tests fail in Intel PROSet.	• Make sure the network cable is plugged into the adapter and link partner. • Check the connector at both ends of the cable. • If the cable is secure, try attaching the cable to a different port on the link partner or try attaching a different cable.
The Connection test in Intel PROSet reports that it	• Verify that your switch is compatible with your network

Problem	Solution
cannot connect at the adapter's maximum speed because the link partner does not support that speed (copper adapters only).	adapter port. • Make sure the link partner is set to auto-negotiate. • For gigabit copper connections, make sure the cable is a 4-pair Category 5 or better. • For 10 gigabit copper connections, make sure the cable is a 4-pair Category 6 or better.
The Cable test fails in Intel PROSet.	Failures of the Cable test indicate breaks, shorts, or mis-wirings in the network cable attached to the adapter. Try reconnecting the cable or using another cable.
The Fan Fail LED of the 10 Gigabit AT Server Adapter is on (red).	The fan cooling solution is not functioning properly. Contact customer support for further instructions.

Other Items to Check

- Use the drivers that came with your adapter, or download the latest ones from Customer Support.
- Make sure the cable is installed properly. The network cable must be securely attached at all connections. If the cable is attached but the problem persists, try a different cable.
- For copper connections, make sure the cable is a 4-pair Category 5 for 1000BASE-T or 100BASE-TX or a 4-pair Category 6 for 10GBASE-T.
- Make sure the link partners match or are set to auto-negotiate. Make sure the updated driver is loaded.
- Test the adapter. Run the adapter and network tests described in Testing the Adapter.
- Check the Common problems table and try the recommended solutions.

Indicator Lights

The Intel® **Ethernet Converged Network Adapter XL710-Q2** has the following indicator lights:

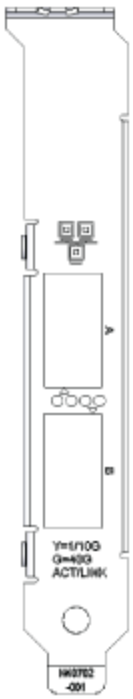

Label	Indication	Meaning
ACT/LNK	Green	Linked at 40 Gb
	Yellow	Linked at 1/10 Gb
	Blinking On/OFF	Actively transmitting or receiving data
	Off	No link.

The **Intel® Ethernet 40G 2P XL710 QSFP+ rNDC** has the following indicator lights:

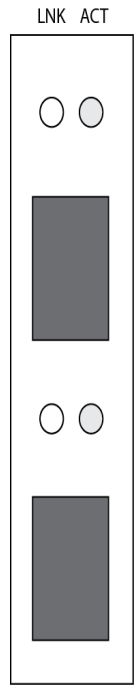

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Ethernet Converged Network Adapter X550-t** has the following indicator lights:

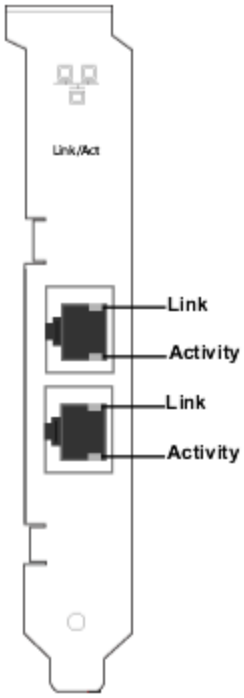

Label	Indication	Meaning
Link	Green	Linked at 10 Gb.
	Yellow	Linked at 1 Gb.
	Off	Linked at 100 Mbps.
Activity	Off	No link.
	Blinking On/Off	Actively transmitting or receiving data.

The **Intel® Ethernet Server Adapter X520-2** has the following indicator lights:

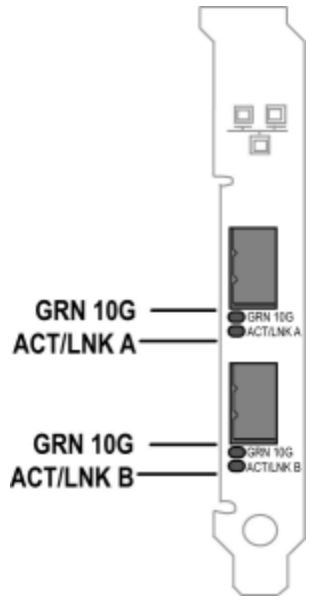

Label	Indication	Meaning
ACT/LNK	Green flashing	Data activity
	Off	No link
10G	Green on	10 Gbps link
	Yellow on	1 Gbps link
	Off	No link

The **Intel® Ethernet Server Adapter X520-T2** has the following indicator lights:

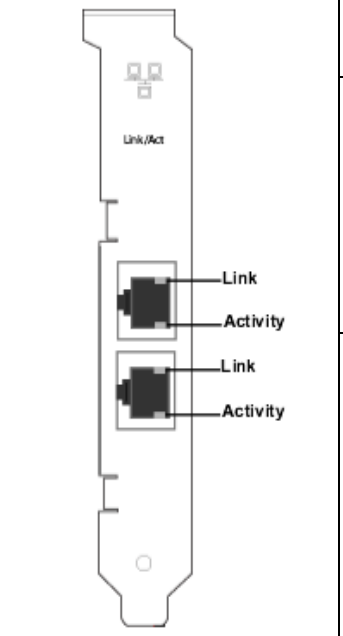

Label	Indication	Meaning
ACT	Green flashing	Data activity
	Off	No link
LNK	Green on	10 Gbps link
	Yellow on	1 Gbps link
	Off	No link

The **Intel® Ethernet 10G 2P X520 Adapters** have the following indicator lights:

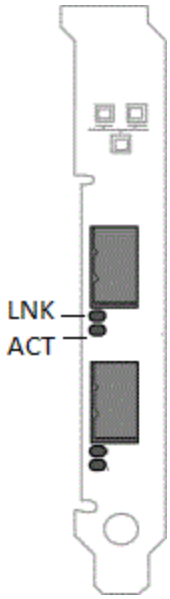

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Ethernet 10G 2P X540-t Adapters** have the following indicator lights:

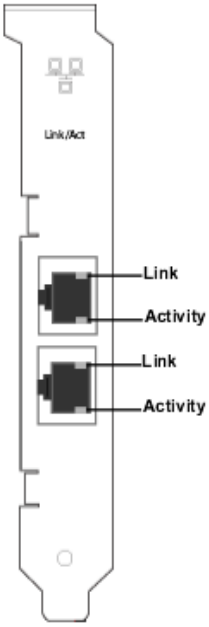

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Ethernet 10G 4P X540/I350 rNDC** and **Intel® Gigabit 4P X540/I350 rNDC** have the following indicator lights:

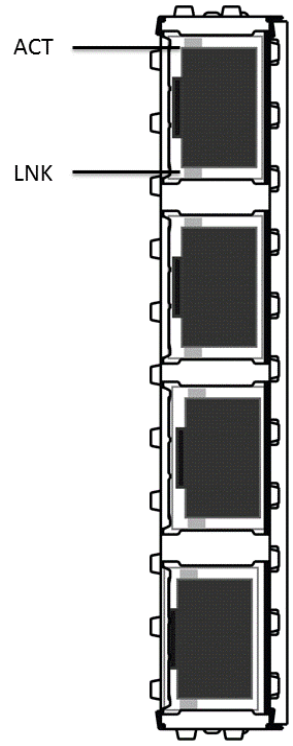

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Ethernet 10G 4P X520/I350 rNDC** and **Intel® Gigabit 4P X520/I350 rNDC** have the following indicator lights:

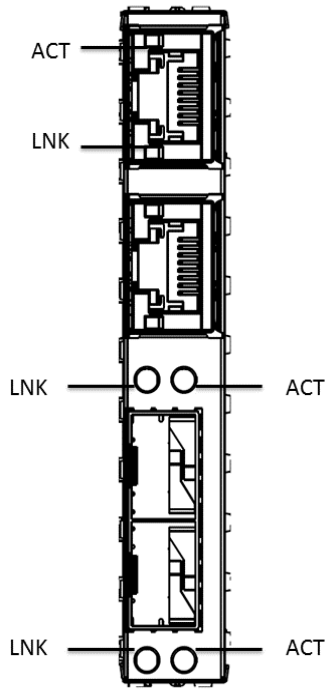

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Gigabit 2P I350-t Adapter** and **Intel® Gigabit 4P I350-t Adapter** and have the following indicator lights:

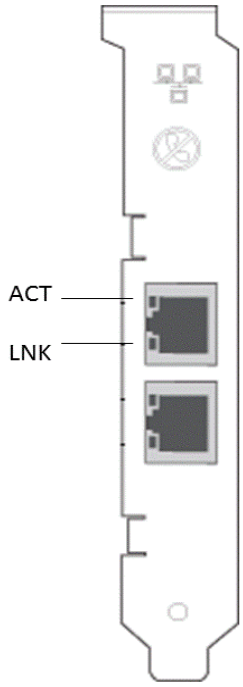

Dual Port

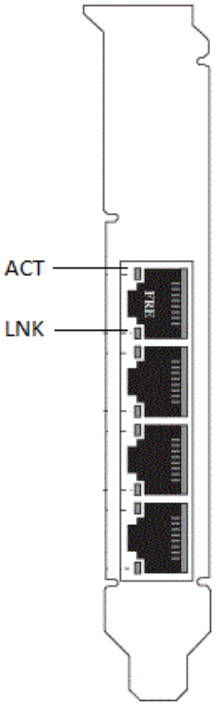

Quad Port

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Gigabit 4P I350-t rNDC** has the following indicator lights:

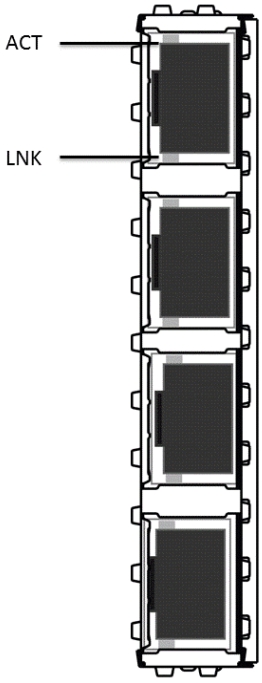

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Ethernet Gigabit 4P x710/I350 rNDC** and **Intel® 10G 4P X710/I350 rNDC** have the following indicator lights:

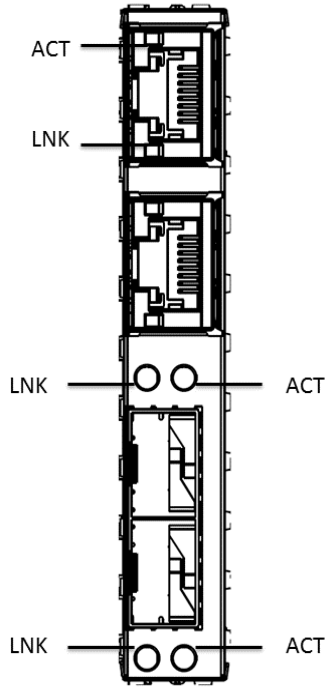

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Converged Network Adapter X710** has the following indicator lights:

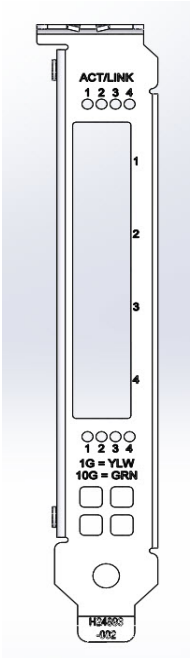

Quad Port

Label	Indication	Meaning
ACT/LINK (green/yellow)	Green Flashing	Data activity on port.
	Yellow on	Operating at lower port speed.
	Off	No link.
1234	Green	Operating at 10G.
	Yellow	Operating at 1G.

The **Intel® Converged Network Adapter X710** has the following indicator lights.

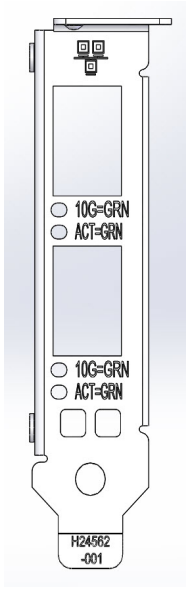

Low Profile

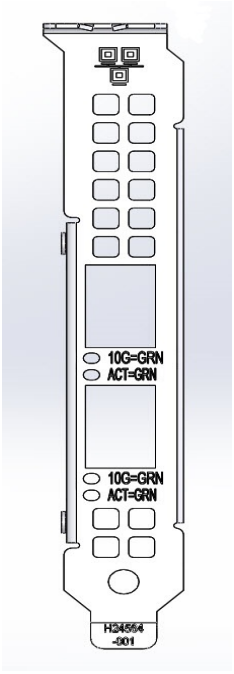

Full Profile

Label	Indication	Meaning
10G (green/yellow)	Green on	Operating at 10G.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

The **Intel® Ethernet 10G 4P x710 SFP+ rNDC** has the following indicator lights:

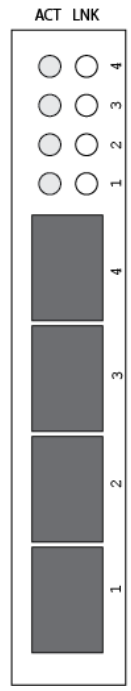

Label	Indication	Meaning
LNK (green/yellow)	Green on	Operating at maximum port speed.
	Yellow on	Operating at lower port speed.
	Off	No link.
ACT (green)	Green flashing	Data activity.
	Off	No activity.

Multiple Adapters

When configuring a multi-adapter environment, you must upgrade all Intel adapters in the computer to the latest drivers and software.

If the computer has trouble detecting all adapters, consider the following:

- Your operating system may need to re-enumerate the bus, especially if any devices with bridge chips are used. To force the re-enumeration, uninstall or unload the drivers for all PCI devices installed and shutdown the computer. Then restart the computer and reinstall or reload all drivers.
- The "Plug and Play OS" setting in the BIOS should be set to "No" for non-Windows operating systems.
- Adapters with the Intel Boot Agent enabled will require a portion of the limited start up memory for each adapter enabled. Disable the service on adapters that do not need to boot Pre-Boot Execution Environment (PXE).

Check the sections on installing driver software for your particular operating system for further information.

Other Performance Issues

Attaining gigabit speeds requires that many components are operating at peak efficiency. Among them are the following:

- **Cable quality and length.** Do not exceed the maximum recommended length for your cable type. Shorter lengths provide better results. Straighten kinks and check for damaged sections of cable.
- **Bus speed and traffic.**
- **Processor speed and load.** Check your performance monitoring programs to see if traffic is being affected by your processor speed, available memory or other processes.
- **Available memory**
- **Transmission frame size.** Your network performance may be enhanced by adjusting or maximizing the transmission frame size. Operating systems, switches and adapters will impose varying limits on maximum frame size. See the discussion on Jumbo Frames.
- **Operating System.** Features vary by OS compatibility, such as offloading and multiprocessor threading.

PCI Express Configuration Troubleshooting

If the adapter is not recognized by your OS or if it does not work you may need to change some BIOS Setup program settings. Try the following only if you are having problems with the adapter.

- Check to see that the "Plug-and-Play" setting is compatible with the operating system you are using.
- Enable the PCI Express slot. In some PCI computers, you may need to use the BIOS Setup program to enable the slot.
- Reserve interrupts and/or memory addresses. This prevents multiple buses or bus slots from using the same interrupts. Check the BIOS for IRQ options.

Known Issues

Code 10 yellow bang errors and devices missing in Windows Device Manager

On an X710 device in a Microsoft Windows Server 2008 R2 system with NParEP enabled, the driver will load on the first 8 physical functions only. The second set of 8 physical functions will have Code 10 yellow bang errors in Windows Device Manager. This is a limitation of the operating system.

Dropped Receive Packets on Half-Duplex 10/100 Networks

If you have an Intel PCI Express adapter installed, running at 10 or 100 Mbps, half-duplex, with TCP Segment Offload (TSO) enabled, you may observe occasional dropped receive packets. To work around this problem, disable TSO or update the network to operate in full-duplex or 1 Gbps.

Throughput Reduction After Hot-Replace

If an Intel gigabit adapter is under extreme stress and is hot-swapped, throughput may significantly drop. This may be due to the PCI property configuration by the Hot-Plug software. If this occurs, throughput can be restored by restarting the system.

CPU Utilization Higher Than Expected

Setting RSS Queues to a value greater than 4 is only advisable for large web servers with several processors. Values greater than 4 may increase CPU utilization to unacceptable levels and have other negative impacts on system performance.

Low-Power Link Speed Slower Than Expected

If you disable the "Reduce Power During Standby" setting and remove power from the system, your system may link at 10 Mbps when power is restored, instead of 100 Mbps or faster. The system will continue to link at 10 Mbps until the operating system is loaded. This setting will be restored when the operating system loads.

VLAN Creation Fails on a Team that Includes a Non-Intel Phantom Adapter

If you are unable to create a VLAN on a team that includes a non-Intel phantom adapter, use Device Manager to remove the team, then recreate the team without the phantom adapter, and add the team to the VLAN.

Team/VLAN Static IP Information May Not Always be Restored

In a Windows Server environment, if you configured team or VLAN static IP information, when you upgrade the Intel Network Connection Software (versions 11.0 or below), the configured items may not always be restored and the default team and VLAN values are used.

A VLAN Created on an Intel Adapter Must be Removed Before a Multi-Vendor Team Can be Created.

In order to create the team, the VLAN must first be removed.

Windows Known Issues

Application Error Event IDs 789, 790, and 791 in the Event Log

If Data Center Bridging (DCB) is enabled, and the enabled port loses link, the following three events may be logged in the event log:

- Event ID 789: Enhanced Transmission Selection feature on a device has changed to non-operational
- Event ID 790: Priority Flow Control feature on a device has changed to non-operational
- Event ID 791: Application feature on a device has changed to non-operational (FCoE)

This is the expected behavior when a DCB enabled port loses link. DCB will begin working again as soon as link is reestablished. A port will lose link if the cable is disconnected, the driver or software package is updated, if the link partner goes down, or for other reasons.

Port is missing from Lifecycle Controller : Network Settings

If a port is configured for iSCSI boot or FCoE boot, and it successfully connected to its boot target, then you cannot modify the port settings in the Lifecycle Controller.

Procedure for Installing and Upgrading Drivers and Utilities

Intel does not recommend installing or upgrading drivers and Intel® PROSet software over a network connection. Instead, install or upgrade drivers and utilities from each system. To install or upgrade drivers and utilities, follow the instructions in the user guide.

"Malicious script detected" Warning from Norton AntiVirus During PROSet Uninstall

The Intel PROSet uninstall process uses a Visual Basic script as part of the process. Norton AntiVirus and other virus scanning software may mistakenly flag this as a malicious or dangerous script. Letting the script run allows the uninstall process to complete normally.

Windows Code 10 Error Message on Driver Install or Upgrade

If you encounter a Windows Code 10 error message when installing or upgrading drivers, reboot to resolve the issue.

Advanced Properties Settings Change While Traffic is Running

In the Advanced Properties tab of Intel® PROSet, parameters should not be modified under heavy network loads. Otherwise, a reboot may be required to make the changes effective.

Unexpected Connectivity Loss

If you uncheck the "Allow the computer to turn off this device to save power" box on the Power Management tab and then put the system to sleep, you may lose connectivity when you exit sleep. You must disable and enable the NIC to resolve the issue. Installing Intel® PROSet for Windows Device Manager will also resolve the issue.

Intermittent Link Loss and Degraded Performance at High Stress Can Occur on Windows Server 2012 Systems

In a Windows Server 2012-based system with multi-core processors, possible intermittent link loss and degraded performance at high stress may occur due to incorrect RSS processor assignments. More information and a Microsoft hotfix are available at: http://support.microsoft.com/kb/2846837.

Virtual machine loses link on a Microsoft Windows Server 2012 R2 system

On a Microsoft Windows Server 2012 R2 system with VMQ enabled, if you change the BaseRssProcessor setting, then install Microsoft Hyper-V and create one or more virtual machines, the virtual machines may lose link. Installing the April 2014 update rollup for Windows RT 8.1, Windows 8.1, and Windows Server 2012 R2 (2919355) and hotfix 3031598 will resolve the issue. See http://support2.microsoft.com/kb/2919355 and http://support2.microsoft.com/kb/3031598 for details.

In a Microsoft Hyper-V environment, Virtual Machines bound to NPAR partitions will not communicate with each other

In a Microsoft Hyper-V environment, if you have NPAR enabled on a port, and Virtual Machines (VMs) bound to partitions on that port, the VMs may not be able to communicate with each other. This happens because the virtual switch inside Hyper-V sends the packets to the physical port, which sends the packets to the switch that is connected to the port. The physical switch may not be configured for reflective relay (also called hairpin mode), so it maynot send the packets back on the same connection from which it received them. Connecting the port to a Virtual Ethernet Port Aggregator (VEPA) capable switch will resolve the issue.

Intel 10GbE Network Adapter Known Issues

The System H/W Inventory (iDRAC) indicates that Auto-negotiation on the Embedded NIC is Disabled, but elsewhere link speed and duplex auto-negotiation is Enabled

If an optical module is plugged into the Intel® Ethernet 10G X520 LOM on a PowerEdge-C6320, the System H/W Inventory (iDRAC) will indicate that Auto-negotiation is Disabled. However, Windows Device Manager and HII indicate that link speed and duplex Auto-negotiation is Enabled. This is because the driver contains an algorithm that allows the LOM to link with SFP partners at 10 Gbps or 1 Gbps. This is reported to Windows Device Manager and HII, but it is not true auto-negotiation. iDRAC reads the device's firmware, which has no knowledge of the algorithm, and therefore reports that auto-negotiation is disabled.

ETS Bandwidth Allocations Don't Match Settings

When Jumbo Frames is set to 9K with a 10GbE adapter, a 90%/10% ETS traffic split will not actually be attained on any particular port, despite settings being made on the DCB switch. When ETS is set to a 90%/10% split, an actual observed split of 70%/30% is more likely.

Supported SFP or SFP+ Module Not Recognized by the System

If you try to install an unsupported module, the port may no longer install any subsequent modules, regardless of whether the module is supported or not. The port will show a yellow bang under Windows Device Manager and an event id 49 (unsupported module) will be added to the system log when this issue occurs. To resolve this issue, the system must be completely powered off.

Link Loss on 10GbE Devices with Jumbo Frames Enabled

You must not lower Receive_Buffers or Transmit_Buffers below 256 if jumbo frames are enabled on an Intel(R) 10GbE Device. Doing so will cause loss of link.

Lower Than Expected Performance on Dual Port 10GbE Devices

Some PCI Express x8 slots are actually configured as x4 slots. These slots have insufficient bandwidth for full 10GbE line rate with dual port 10GbE devices. The driver can detect this situation and will write the following message in the system log: "PCI Express bandwidth available for this card is not sufficient for optimal performance. For optimal performance a x8 PCI Express slot is required." If this error occurs, moving your adapter to a true x8 slot will resolve the issue.

Failed connection and possible system instability

If you have non-Intel networking devices capable of Receive Side Scaling installed in your system, the Microsoft Windows registry keyword “RSSBaseCPU” may have been changed from the default value of 0x0 to point to a logical processor. If this keyword has been changed, then devices based on Intel(R) 82598 or 82599 10 Gigabit Ethernet Controllers might not pass traffic. Attempting to make driver changes in this state may cause system instability. Set the value of RSSBaseCpu to 0x0, or to a value corresponding to a physical processor, and reboot the system to resolve the issue.

Intel® Ethernet 10G 2P/4P X710-k bNDC does not have link and is not displayed in Windows Device Manager

If you install an Intel® Ethernet 10G 2P X710-k bNDC or an Intel® Ethernet 10G 4P X710-k bNDC onto a Dell PowerEdge M630/M830 blade server, and install that blade into an M1000e chassis, the bNDC may not have link and may display a yellow bang, or may not be displayed at all, in Windows Device Manager. This is limited to the 1.0 version of the M1000e Midplane.

Intel® Ethernet 10G X520 LOM links at 10 Gbps when 1.0 Gbps Full Duplex is selected

When connected with a direct attach cable, the Intel® Ethernet 10G X520 LOM will always connect at 10 Gbps.

Windows* Event Log

Windows Event Log Service Names

Intel® Ethernet Controller	NDIS Driver File Names	Windows Event Log Service Name
I350	E1r*.sys	e1rexpress
I354	E1s*.sys	e1sexpress
X520	Ixn*.sys	ixgbn
X540	Ixt*.sys	ixgbt
X550	Ixs*.sys	ixgbs
X710	I40ea*.sys	i40ea

Intel® Network Adapter Messages

Below is a list of custom event messages that appear in the Windows Event Log for Intel® Ethernet adapters:

Event ID	Message	Severity
6	PROBLEM: Unable to allocate the map registers necessary for operation. ACTION: Reduce the number of transmit descriptors and restart.	Error
7	PROBLEM: Could not assign an interrupt for the network adapter. ACTION: Try a different PCIe slot. ACTION: Install the latest driver from http://www.intel.com/support/go/network/adapter/home.htm.	Error
23	PROBLEM: The EEPROM on the network adapter may be corrupt. ACTION: Visit the support web site at http://www.intel.com/support/go/network/adapter/home.htm.	Error
24	PROBLEM: Unable to start the network adapter. ACTION: Install the latest driver from http://www.intel.com/support/go/network/adapter/home.htm.	Error
25	PROBLEM: The MAC address on the network adapter is invalid. ACTION: Visit http://www.intel.com/support/go/network/adapter/home.htm for assistance.	Error
27	Network link has been disconnected.	Warning
30	PROBLEM: The network adapter is configured for auto-negotiation but the link partner is not. This may result in a duplex mismatch. ACTION: Configure the link partner for auto-negotiation.	Warning
31	Network link has been established at 10 Gbps full duplex.	Informational
32	Network link has been established at 1 Gbps full duplex.	Informational
33	Network link has been established at 100 Mbps full duplex.	Informational
34	Network link has been established at 100 Mbps half duplex.	Informational
35	Network link has been established at 10 Mbps full duplex.	Informational
36	Network link has been established at 10 Mbps half duplex.	Informational
37	PROBLEM: PCI Express bandwidth available for this adapter is not sufficient for optimal performance. ACTION: Move the adapter to a x8 PCI Express slot.	Warning
40	Intel Smart Speed has downgraded the link speed from the maximum advertised.	Informational
41	The network adapter driver has been stopped.	Informational
42	The network adapter driver has been started.	Informational
43	PROBLEM: Could not allocate shared memory necessary for operation. ACTION: Reduce the number of transmit and receive descriptors, then restart.	Error
44	PROBLEM: Could not allocate memory necessary for operation. ACTION: Reduce the number of transmit and receive descriptors, then restart.	Error
45	PROBLEM: Could not allocate a resource pool necessary for operation. ACTION: Reduce the number of transmit and receive descriptors, then restart.	Error
46	PROBLEM: Could not initialize scatter-gather DMA resources necessary for operation. ACTION: Reduce the number of transmit descriptors and restart.	Error
47	PROBLEM: Could not map the network adapter flash. ACTION: Install the latest driver from http://www.intel.com/support/go/network/adapter/home.htm.	Error

Event ID	Message	Severity
	tel.com/support/go/network/adapter/home.htm. ACTION: Try another slot.	
48	PROBLEM: The fan on the network adapter has failed. ACTION: Power off the machine and replace the network adapter.	Error
49	PROBLEM: The driver was unable to load due to an unsupported SFP+ module installed in the adapter. ACTION: Replace the module. ACTION: Install the latest driver from http://www.in-tel.com/support/go/network/adapter/home.htm.	Error
50	PROBLEM: The network adapter has been stopped because it has overheated. ACTION: Restart the computer. If the problem persists, power off the computer and replace the network adapter.	Error
51	PROBLEM: The network adapter link speed was downshifted because it overheated.	Error
52	PROBLEM: The network adapter has been stopped because it has overheated.	Error
53	Jumbo Frames cannot be configured when MACSec is enabled.	Informational
54	PROBLEM: A malicious VF driver has been detected.	Warning
56	The network driver has been stopped because the network adapter has been removed.	Informational
58	Network link has been established at 25Gbps full duplex.	Informational
60	Network link has been established at 50Gbps full duplex.	Informational
61	Network link has been established at 20Gbps full duplex.	Informational
64	This network adapter's etrack ID is:	Informational
65	PROBLEM: PCI Express bandwidth available for this adapter is not sufficient for optimal performance. ACTION: Move the adapter to a Generation 3 x4 PCI Express slot.	Warning
66	PROBLEM: PCI Express bandwidth available for this adapter is not sufficient for optimal performance. ACTION: Move the adapter to a Generation 3 x8 PCI Express slot.	Warning
67	The partition detected link speed that is less than 10Gbps.	Warning
68	The driver for the device stopped because the NVM image is newer than the driver. You must install the most recent version of the network driver.	Error
69	The driver for the device detected a newer version of the NVM image than expected. Please install the most recent version of the network driver.	Warning
70	The driver for the device detected an older version of the NVM image than expected. Please update the NVM image.	Informational
71	The driver failed to load because an unsupported module type was detected.	Error
72	PROBLEM: The driver failed to load because the adapter was not provided MSI-X interrupt resources. ACTION: Move the adapter to another slot or platform.	Error
73	The 'Speed and Duplex' and 'Flow Control' user settings cannot be changed since this device is operating in virtual connect mode.	Informational

Intel Advanced Network Services Messages

Below is a list of intermediate driver custom event messages that appear in the Windows Event Log:

Event ID	Message	Severity
2	Unable to allocate required resources. Free some memory resources and restart.	Error
3	Unable to read required registry parameters. To resolve, remove the adapter team and then create a new team.	Error
4	Unable to bind to physical adapter. To resolve, remove the adapter team and then create a new team.	Error
5	Unable to initialize an adapter team. To resolve, remove the adapter team and then create a new team.	Error
6	Primary Adapter is initialized: <member description>	Informational
7	Adapter is initialized: <member description>	Informational
8	Team #<team ID>: Team is initialized.	Informational
9	Team #<ID>: Virtual Adapter for <VLAN name> [VID=<VLAN ID>] initialized.	Informational
10	Current Primary Adapter is switching from: <member description>	Informational
11	Adapter link down: <member description>	Warning
12	Secondary Adapter took over: <member description>	Informational
13	The <member description> has been deactivated from the team.	Warning
14	Secondary Adapter has rejoined the Team: <member description>	Informational
15	Adapter link up: <member description>	Informational
16	Team #<ID>: The last adapter has lost link. Network connection has been lost.	Error
17	Team #<ID>: An adapter has re-established link. Network connection has been restored.	Informational
18	Preferred primary adapter has been detected: <member description>	Informational
19	Preferred secondary adapter has been detected: <member description>	Informational
20	Preferred primary adapter took over: <member description>	Informational
21	Preferred secondary adapter took over: <member description>	Informational
22	Primary Adapter does not sense any Probes: <member description>. Possible reason: partitioned Team.	Warning
23	Team #<ID>: A Virtual Adapter failed to initialize.	Error
32	An illegal loopback situation has occurred on the adapter in device <member description>. Check the configuration to verify that all the adapters in the team are connected to 802.3ad compliant switch ports.	Warning
35	Initializing Team #<ID> with <missing #> missing adapters. Check the configuration to verify that all the adapters are present and functioning.	Warning
37	Virtual adapter for <VLAN name> [VID=<VLAN ID>] removed from team #<team ID>.	Informational
38	Adapter removed from team #<ID>.	Informational
39	You may not be able to change the virtual adapter settings. To resolve, reload the driver.	Warning
40	Virtual adapter unload process may have not completed successfully. Driver may not be unloaded. To resolve, reboot the system.	Warning

Intel DCB Messages

Below is a list of intermediate driver custom event messages that appear in the Windows Event Log:

Event ID	Message	Severity
256	Service debug string	Informational
257	Enhanced Transmission Selection feature has been enabled on a device.	Informational
258	Enhanced Transmission Selection feature has been disabled on a device.	Informational
259	Priority Flow Control feature has been enabled on a device.	Informational
260	Priority Flow Control feature has been disabled on a device.	Informational
261	Enhanced Transmission Selection feature on a device has changed to operational.	Informational
262	Priority Flow Control feature on a device has changed to operational.	Informational
263	Application feature on a device has changed to operational.	Informational
264	Application feature has been disabled on a device.	Informational
265	Application feature has been enabled on a device.	Informational
269	Logical Link feature on a device has changed to operational.	Informational
270	Logical Link feature has been disabled on a device.	Informational
271	Logical Link feature has been enabled on a device.	Informational
768	Service failed while starting.	Error
770	Service handler failed while installing.	Error
771	Service could not allocate sufficient memory.	Error
772	Service unable to use network adapter.	Error
773	Service rejected configuration - invalid total for transmit bandwidth groups.	Error
774	Service rejected configuration - invalid total for receive bandwidth groups.	Error
775	Service rejected configuration - invalid transmit bandwidth group index.	Error
776	Service rejected configuration - invalid receive bandwidth group index.	Error
777	Service rejected configuration - link strict and non-zero bandwidth on transmit traffic class.	Error
778	Service rejected configuration - link strict and non-zero bandwidth on receive traffic class.	Error
779	Service rejected configuration - zero bandwidth on transmit traffic class.	Error
780	Service rejected configuration - zero bandwidth on receive traffic class.	Error
781	Service rejected configuration - link strict and non-zero bandwidth on transmit bandwidth group.	Error
782	Service rejected configuration - link strict and non-zero bandwidth on receive bandwidth group.	Error
783	Service rejected configuration - invalid total transmit for bandwidth group.	Error
784	Service rejected configuration - invalid total receive for bandwidth group.	Error
785	Service unable to configure needed WMI services.	Error
786	Service experienced a transmit state machine error.	Error

Event ID	Message	Severity
787	Service experienced a receive state machine error.	Error
789	Service connection to LLDP protocol driver failed.	Error
790	Enhanced Transmission Selection feature on a device has changed to non-operational.	Error
791	Priority Flow Control feature on a device has changed to non-operational.	Error
792	Application feature on a device has changed to non-operational.	Error
793	Service rejected configuration - multiple link strict bandwidth groups were detected.	Error
794	Logical Link feature on a device has changed to non-operational.	Error
795	Failed to open device.	Error
796	DCB settings of the network adapter are invalid.	Error
797	DCB settings of the network adapter are invalid - AppSelector.	Error
798	Detected a non-optimal network adapter driver component. Please install network adapter driver version 3.5 or greater.	Error

Intel iSCSI DCB Messages

Below is a list of intermediate driver custom event messages that appear in the Windows Event Log:

Event ID	Message	Severity
4352	Service debug string:	Informational
4353	iSCSI DCB Agent has added a QOS filter for iSCSI traffic.	Informational
4354	iSCSI DCB Agent has removed a QOS filter for iSCSI traffic.	Informational
4355	iSCSI DCB Agent has modified a QOS filter for iSCSI traffic.	Informational
4356	iSCSI DCB Agent was notified by the QOS service that an iSCSI DCB adapter was closed.	Informational
4357	Priority Flow Control and Application User Priority are configured for iSCSI DCB traffic.	Informational
4358	All members of the Team configured for iSCSI DCB traffic have a valid DCB configuration.	Informational
8704	Some members of the Team configured for iSCSI DCB traffic have an invalid DCB configuration.	Warning
13056	Service failed while starting.	Error
13057	Service handler failed while installing.	Error
13058	Error returned by Traffic Control interface.	Error
13059	Service could not allocate sufficient memory.	Error
13060	iSCSI DCB Agent is unable to add the QOS filter for iSCSI traffic.	Error
13061	iSCSI DCB Agent was notified by the QOS service that all QOS filters for an iSCSI DCB adapter were removed.	Error
13062	Application User Priority or Priority Flow Control is misconfigured for iSCSI DCB traffic.	Error
13063	Priority Flow Control TLV is non-operational for iSCSI DCB traffic.	Error
13064	Application TLV is non-operational for iSCSI DCB traffic.	Error

Event ID	Message	Severity
13065	Detected unsupported Operating System.	Error
13066	No member of the Team configured for iSCSI DCB traffic has a valid DCB configuration.	Error

Adapter Test Utilities

Overview

Intel's diagnostic software lets you test the adapter to see if there are any problems with the adapter hardware, the cabling, or the network connection. You can also use diagnostics to isolate problems during troubleshooting.

DIAGS.EXE runs under MS-DOS* and later compatible operating systems. It will not run from a Windows* Command Prompt within any version of the Microsoft Windows operating system or in any other non-MS-DOS operating system.

This utility is designed to test hardware operation and confirm the adapter's ability to communicate with another adapter in the same network. It is not a throughput measurement tool.

DIAGS can test the adapter whether or not there is a responder present. In order to do a thorough test, however, you should set up a second system on the network as a responder prior to starting a test. If there are hot keys, the letters will be highlighted.

Starting the Test Utility

NOTE: If there is an MS-DOS network driver present, such as NDIS2 or DOS-ODI, the test utility and the network driver could become unstable. You should reboot and ensure that there are no network drivers loaded.

1. Boot to MS-DOS.
2. Navigate to the \DOSUtilities\UserDiag directory, then type `DIAGS` at the prompt and press <Enter>.

   The test utility program automatically scans the hardware and lists all Intel-based adapters. They are listed in this manner:
   - If you have only one network connection in your system, this screen will be bypassed.
   - If you have a dual-port or quad-port adapter, each port is listed separately starting with Port A, then Port B, etc. You can find the port information on the bracket label.

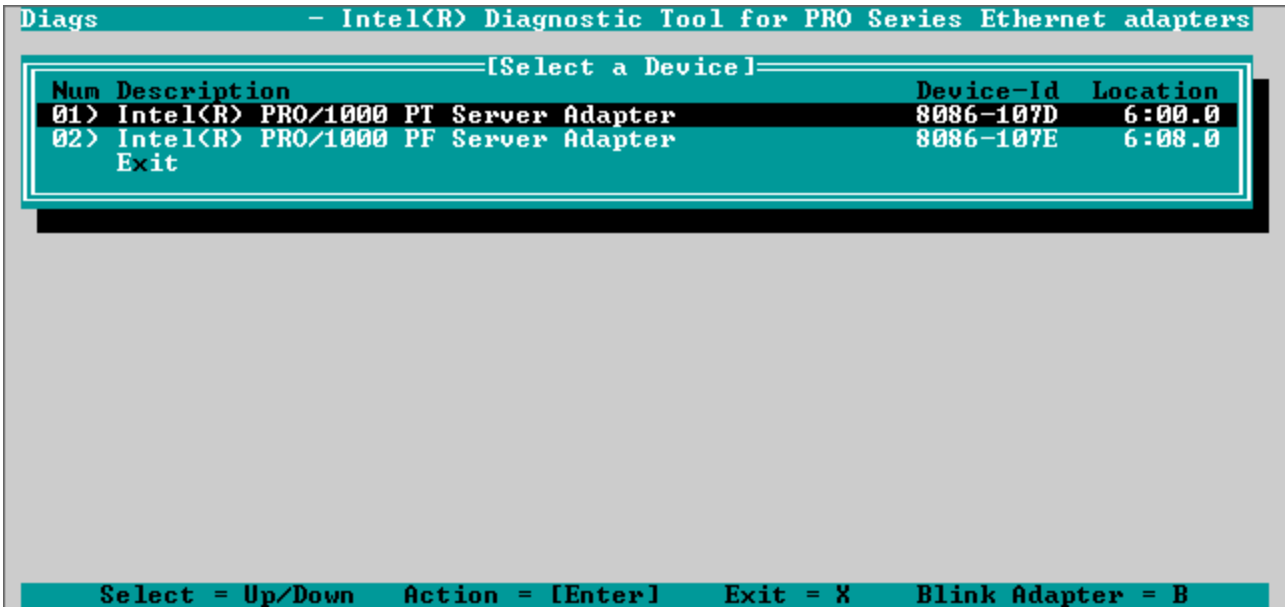

3. Select the adapter you want to test by moving the highlight and pressing <Enter>. The test utility program displays its main menu.

View Adapter Configuration

Selecting **View Adapter Configuration** will bring up the adapter configuration screen. This screen describes various properties of the adapter.

Press <F5> to view additional information on the PCI Express slot occupied by the adapter. This information is primarily used for troubleshooting by Customer Support.

Press any key to return to Adapter Configuration.

Test Adapter Menu

Selecting **Test Adapter** from the Main Menu brings up the Test Menu. This menu allows you to select which tests to perform on the adapter and configure the test options.

Begin Adapter Tests

Selecting this option brings up the test screen. While tests are being performed, a rotating spinner is shown to indicate the application is still "alive." The results of the tests are displayed as each test is performed. If multiple test passes are selected, the results contain a count of test failures. A list containing zeros means that all tests have passed. Single tests will display "Passed" or "Failed" for each pass.

Change Test Options

The test setup screen allows you to configure and select the specific tests desired. Each option is toggled by moving the cursor with the arrow keys and pressing <Enter> to change the option. The number of tests is simply entered from the keyboard in the appropriate box. If there is a gap in the menu, that means the test is not supported by your adapter. By default, local diagnostics run automatically, while network diagnostics are disabled.

NOTE: The test program will test attributes that are applicable to your adapter. Only supported tests are displayed.

Device Registers - Test patterns are written, read, and verified through the adapter's device registers to ensure proper functionality.

FIFOs - This will write test bit patterns to the adapter's FIFO buffers to make sure the FIFOs are working properly. Not all adapters have FIFO, so it will not appear in all test lists.

EEPROM - This test tests both the readability of the EEPROM as well as the integrity of the data stored in the EEPROM. It reads the EEPROM and calculates the checksum. This checksum is then compared to the checksum stored in the EEPROM. If values are not the same, the test reports failure.

Interrupt - This tests the adapter's ability to generate an interrupt and have it propagated through the system to the Programmable Interrupt Controller (PIC). The test triggers an interrupt by setting the interrupt cause register and then verifies that an interrupt has been triggered.

Loopback - There are two internal loopback tests. These tests set the adapter in the appropriate loopback mode and send packets back through the adapter's receive circuitry and logic. These tests are chipset-dependent and may not be selectable.

Link - Checks to see if the adapter has link or does not have link.

Network Test - The Network Test looks for a responder, and then sends packets. If no responder is found, the test reports failure. If packets are received back from the responder, the test reports success.

NOTE: In some instances, the test may fail when it is connected to a switch with Spanning Tree Protocol enabled.

Networking Menu

The networking menu contains network-specific tests, such as Spanning Tree detection and Network test responder.

Set Up as Responder

This allows the user to set up the adapter as a responder so a connected system can perform the network test portion of the diagnostics tests. Although you can use a variety of adapters as the responder and connect directly or through a switch, the best results are obtained with a cross-over cable and a same-type adapter.

When you press <Esc>, the responder operation is canceled and control is immediately returned to the Networking menu.

Detect Spanning Tree

Spanning trees can be troublesome in a networking configuration. The Detect Spanning Tree option attempts to detect if a spanning tree exists on the network. This is done by resetting the link and listening for spanning tree packets.

Regulatory Compliance Statements

FCC Class A Products

40 Gigabit Ethernet Products

- Intel® Ethernet 40G 2P XL710 QSFP+ rNDC
- Intel® Ethernet Converged Network Adapter XL710-Q2

10 Gigabit Ethernet Products

- Intel® Ethernet X520 10GbE Dual Port KX4-KR Mezz
- Intel® Ethernet 10G 2P X540-t Adapter
- Intel® Ethernet 10G 2P X550-t Adapter
- Intel® Ethernet 10G 4P X540/I350 rNDC
- Intel® Ethernet 10G 4P X520/I350 rNDC
- Intel® Ethernet 10G 2P X520-k bNDC
- Intel® Ethernet 10G 4P X710-k bNDC
- Intel® Ethernet 10G 2P X710-k bNDC
- Intel® Ethernet 10G X710-k bNDC
- Intel® Converged Network Adapter X710
- Intel® Ethernet 10G 4P X710/I350 rNDC
- Intel® Ethernet 10G 4P X710 SFP+ rNDC
- Intel® Ethernet 10G X710 rNDC

Gigabit Ethernet Products

- Intel® Gigabit 4P I350-t rNDC
- Intel® Gigabit 4P X540/I350 rNDC
- Intel® Gigabit 4P X520/I350 rNDC
- Intel® Gigabit 4P I350-t Mezz
- Intel® Gigabit 4P X710/I350 rNDC
- Intel® Gigabit 4P I350-t bNDC

FCC Class B Products

10 Gigabit Ethernet Products

- Intel® Ethernet 10G 2P X520 Adapter
- Intel® Ethernet 10G X520 LOM

Gigabit Ethernet Products

- Intel® Gigabit 2P I350-t Adapter
- Intel® Gigabit 4P I350-t Adapter
- Intel® Gigabit 4P I350 bNDC

Safety Compliance

The following safety standards apply to all products listed above.

- UL 60950-1, 2nd Edition, 2011-12-19 (Information Technology Equipment - Safety - Part 1: General Requirements)
- CSA C22.2 No. 60950-1-07, 2nd Edition, 2011-12 (Information Technology Equipment - Safety - Part 1: General Requirements)
- EN 60950-1:2006/A11:2009/A1:2010/A12:2011 (European Union)
- IEC 60950-1:2005 (2nd Edition); Am 1:2009 (International)
- EU LVD Directive 2006/95/EC

EMC Compliance – The following standards may apply:

Class A products:

- FCC Part 15 – Radiated & Conducted Emissions (USA)
- CAN ICES-3(A)/NMB-3(A) – Radiated & Conducted Emissions (Canada)
- CISPR 22 – Radiated & Conducted Emissions (International)
- EN55022: 2010 – Radiated & Conducted Emissions (European Union)
- EN55024: 2010 +A1:2001+A2:2003 – Immunity (European Union)
- EMC Directive 2004/108/EC
- VCCI (Class A)– Radiated & Conducted Emissions (Japan)
- CNS13438 – Radiated & Conducted Emissions (Taiwan)
- AS/NZS CISPR 22 – Radiated & Conducted Emissions (Australia/New Zealand)
- NRRA No. 2012-13 (2012.06.28), NRRA Notice No. 2012-14 (2012.06.28) (Korea)

Class B products:

- FCC Part 15 (Class B) – Radiated & Conducted Emissions (USA)
- CAN ICES-3(B)/NMB-3(B) – Radiated & Conducted Emissions (Canada)
- CISPR 22 – Radiated & Conducted Emissions (International)
- EN55022: 2010 – Radiated & Conducted Emissions (European Union)
- EN55024: 2010 – Immunity (European Union)
- EU – EMC Directive 2004/108/EC
- VCCI (Class B)– Radiated & Conducted Emissions (Japan) (excluding optics)
- CNS13438 (Class B)-2006 – Radiated & Conducted Emissions (Taiwan) (excluding optics)
- AS/NZS CISPR 22 – Radiated & Conducted Emissions (Australia/New Zealand)
- KN22; KN24 – Korean emissions and immunity
- NRRA No. 2012-13 (2012.06.28), NRRA Notice No. 2012-14 (2012.06.28) (Korea)

Regulatory Compliance Markings

When required these products are provided with the following Product Certification Markings:

- UL Recognition Mark for USA and Canada
- CE Mark
- EU WEEE Logo
- FCC markings
- VCCI marking
- Australian C-Tick Mark
- Korea MSIP mark
- Taiwan BSMI mark
- People's Republic of China "EFUP" mark

FCC Class A User Information

The Class A products listed above comply with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

CAUTION: If the device is changed or modified without permission from Intel, the user may void his or her

authority to operate the equipment.

Canadian Compliance (Industry Canada)

CAN ICES-3(A)/NMB-3(A)

VCCI Class A Statement

この装置は、クラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 ＶＣＣＩ－Ａ

BSMI Class A Statement

警告使用者:
此為甲類資訊技術設備，於居住環境中使用時，
可能會造成射頻擾動，在此種情況下，使用者會
被要求採取某些適當的對策。

KCC Notice Class A (Republic of Korea Only)

A급 기기 (업무용 방송통신기기) **CLASS A device** (commercial broadcasting and communication equipment)	이 기기는 업무용(A급)으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다. This device has been approved by EMC registration. Distributors or users pay attention to this point. This device is usually aimed to be used in other area except at home.

BSMI Class A Notice (Taiwan)

警告使用者:
此為甲類資訊技術設備，於居住環境中使用時，
可能會造成射頻擾動，在此種情況下，使用者會
被要求採取某些適當的對策。

FCC Class B User Information

This equipment has been tested and found to comply with the limits for a Class B digital device pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

CAUTION: If the device is changed or modified without permission from Intel, the user may void his or her authority to operate the equipment.

NOTE: This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Electromagnetic Compatibility Notices

FCC Declaration of Conformity Statement

The following products have been tested to Comply with FCC Standards for Home or Office Use.

PRO/1000 MT, PRO/1000 PT, PRO/1000 GT, Gigabit PT, Gigabit ET, I210-T1, I340-T2/T4, I350-T2/T4, PRO/100 M Desktop Adapter, PRO/100 S Desktop Adapter, PRO/100 S Server Adapter, and PRO/100 S Dual Port Server Adapter

Canadian Compliance (Industry Canada)

CAN ICES-3 (B)/NMB-3 (B)

VCCI Class B Statement (Japan)

この装置は、クラスＢ情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。

ＶＣＣＩ－Ｂ

KCC Notice Class B (Republic of Korea Only)

B급 기기 (가정용 방송통신기기) **CLASS B device** residential broadcasting and communication equipment	이 기기는 가정용(B급)으로 전자파적합등록을 한 기기로서 주로 가정에서 사용하는 것을 목적으로 하며, 모든 지역에서 사용할 수 있습니다. This device has been approved by EMC Registration and is usually aimed to be used in a residential area so that it can be used in all other location as well as at home.

EU WEEE Logo

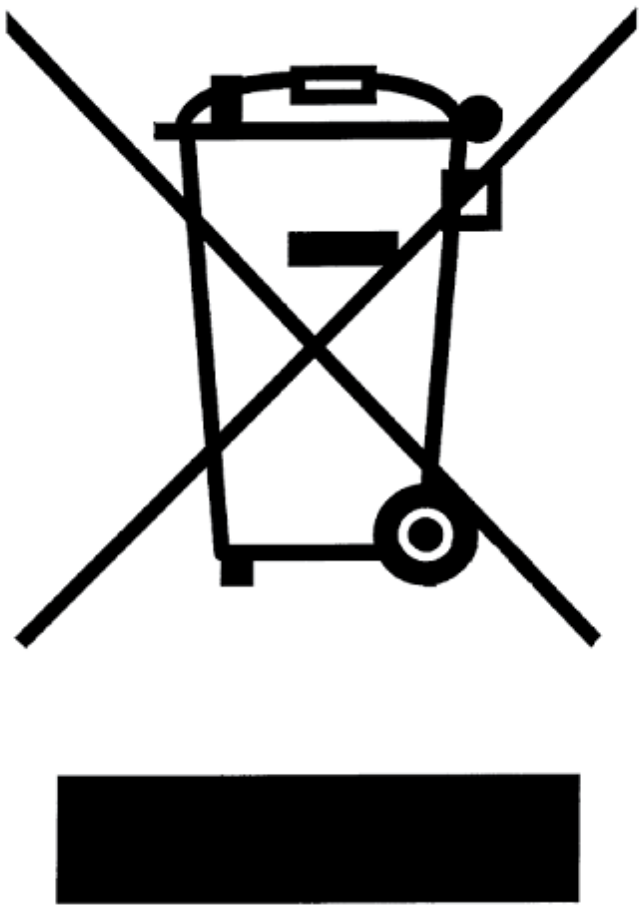

Manufacturer Declaration European Community

Manufacturer Declaration

Intel Corporation declares that the equipment described in this document is in conformance with the requirements of the European Council Directive listed below:

- Low Voltage Directive 2006/95/EC
- EMC Directive2004/108/EC
- RoHS Directive 2011/65/EU

These products follow the provisions of the European Directive 1999/5/EC.

Dette produkt er i overensstemmelse med det europæiske direktiv 1999/5/EC.

Dit product is in navolging van de bepalingen van Europees Directief 1999/5/EC.

Tämä tuote noudattaa EU-direktiivin 1999/5/EC määräyksiä.

Ce produit est conforme aux exigences de la Directive Européenne 1999/5/EC.

Dieses Produkt entspricht den Bestimmungen der Europäischen Richtlinie 1999/5/EC.

Þessi vara stenst reglugerð Evrópska Efnahags Bandalagsins númer 1999/5/EC.

Questo prodotto è conforme alla Direttiva Europea 1999/5/EC.

Dette produktet er i henhold til bestemmelsene i det europeiske direktivet 1999/5/EC.

Este produto cumpre com as normas da Diretiva Européia 1999/5/EC.

Este producto cumple con las normas del Directivo Europeo 1999/5/EC.

Denna produkt har tillverkats i enlighet med EG-direktiv 1999/5/EC.

This declaration is based upon compliance of the Class A products listed above to the following standards:
EN 55022:2010 (CISPR 22 Class A) RF Emissions Control.
EN 55024:2010 (CISPR 24) Immunity to Electromagnetic Disturbance.
EN 60950-1:2006/A11:2009A1:2010/A12:2011 Information Technology Equipment- Safety-Part 1: General Requirements.
EN 50581:2012 - Technical documentation for the assessment of electrical and electronic products with respect to the restriction of hazardous substances.

This declaration is based upon compliance of the Class B products listed above to the following standards:
EN 55022:2010 (CISPR 22 Class B) RF Emissions Control.
EN 55024:2010 (CISPR 24) Immunity to Electromagnetic Disturbance.
EN 60950-1:2006/A11:2009/A1:2010/A12:2011 Information Technology Equipment- Safety-Part 1: General Requirements.
EN 50581:2012 - Technical documentation for the assessment of electrical and electronic products with respect to the restriction of hazardous substances.

WARNING: In a domestic environment, Class A products may cause radio interference, in which case the user may be required to take adequate measures.

Responsible Party

Intel Corporation, Mailstop JF3-446
5200 N.E. Elam Young Parkway
Hillsboro, OR 97124-6497
Phone 1-800-628-8686

China RoHS Declaration

关于符合中国《电子信息产品污染控制管理办法》的声明
Management Methods on Control of Pollution From
Electronic Information Products
(China RoHS declaration)

产品中有毒有害物质的名称及含量

部件名称	有毒有害物质或元素					
	铅 (Pb)	汞 (Hg)	镉 (Cd)	六价铬 (Cr(VI))	多溴联苯 (PBB)	多溴二苯醚 (PBDE)
印刷板组件	X	O	O	O	O	O

O：表示该有毒有害物质在该部件所有均质材料中的含量均在 SJ/T 11363-2006 标准规定的限量要求以下。

X：表示该有毒有害物质至少在该部件的某一均质材料中的含量超出 SJ/T 11363-2006 标准规定的限量要求。

Class 1 Laser Products

Server adapters listed above may contain laser devices for communication use. These devices are compliant with the requirements for Class 1 Laser Products and are safe in the intended use. In normal operation the output of these laser devices does not exceed the exposure limit of the eye and cannot cause harm.

For continued safe operation in case of an abnormal circumstance, always have the provided laser connector cover in place or a compatible fiber optics cable properly connected when power is available to the product.

The Laser device must be factory serviced ONLY by the responsible manufacturer! NO adjustments, service or maintenance is to be performed otherwise.

CAUTION: Use of controls or adjustments or performance of procedures other than those specified herein may result in hazardous radiation exposure.

These Class 1 Laser devices:

Comply with FDA/CDRH per CFR21, subchapter J.
Comply with IEC 60825-1:2007

End-of-Life / Product Recycling

Product recycling and end-of-life take-back systems and requirements vary by country.

Contact the retailer or distributor of this product for information about product recycling and/or take-back.

Support

Web and Internet Sites

http://support.dell.com/

Customer Support Technicians

If the troubleshooting procedures in this document do not resolve the problem, please contact Dell, Inc. for technical assistance (refer to the "Getting Help" section in your system documentation).

Before you call...

You need to be at your computer with your software running and the product documentation at hand.

The technician may ask for the following:

- Your address and telephone number
- The name and model number of the product you are calling about
- The serial number and service tag of the product
- The names and version numbers of the software you are using to operate the product
- The name and version number of the operating system you are using
- The computer type (manufacturer and model number)
- Expansion boards or add-in cards in your computer
- The amount of memory in your computer

Adapter Specifications

Intel® 40 Gigabit Network Adapter Specifications

Feature	Intel® Ethernet Converged Network Adapter XL710-Q2
Bus Connector	PCI Express 3.0
Bus Speed	x8
Transmission Mode/Connector	QSFP+
Cabling	40GBase-SR4, Twinax DAC (7m max)
Power Requirements	6.5 W Maximum @ +12 V
Dimensions (excluding bracket)	5.21 x 2.71 in 13.3 x 6.9 cm
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)
MTBF	159 years
Available Speeds	10 Gbps/40 Gbps
Duplex Modes	Full only
Indicator Lights	*Two per port:* Link and Activity
Standards Conformance	IEEE 802.3ba SFF-8436 PCI Express 3.0
Regulatory and Safety	**Safety Compliance** ● UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) ● EN 60 950 (European Union) ● IEC 60 950 (International) **EMC Compliance** ● FCC Part 15 - Radiated & Conducted Emissions (USA) ● ICES-003 - Radiated & Conducted Emissions (Canada) ● CISPR 22 - Radiated & Conducted Emissions (International) ● EN55022-1998 - Radiated & Conducted Emissions (European Union) ● EN55024 - 1998 - (Immunity) (European Union) ● CE - EMC Directive (89/336/EEC) (European Union) ● VCCI - Radiated & Conducted Emissions (Japan) ● CNS13438 - Radiated & Conducted Emissions (Taiwan) ● AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) ● MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)

Intel® 40GbE Network Daughter Cards (NDC) Specifications

Feature	Intel® Ethernet 40G 2P XL710 QSFP+ rNDC
Bus Connector	PCI Express 3.0
Bus Speed	x8

Transmission Mode/Connector	QSFP+
Cabling	40GBase-SR4, Twinax DAC (7m max)
Power Requirements	6.2 W Maximum @ +12 V
Dimensions (excluding bracket)	3.66 x6.081 in 9.3 x 15.5 cm
Operating Temperature	32 - 140 deg. F (0 - 60 deg. C)
MTBF	112 years
Available Speeds	10 Gbps/40 Gbps
Duplex Modes	Full only
Indicator Lights	*Two per port:* Link and Activity
Standards Conformance	IEEE 802.3ba SFF-8436 PCI Express 3.0
Regulatory and Safety	**Safety Compliance** • UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) • EN 60 950 (European Union) • IEC 60 950 (International) **EMC Compliance** • FCC Part 15 - Radiated & Conducted Emissions (USA) • ICES-003 - Radiated & Conducted Emissions (Canada) • CISPR 22 - Radiated & Conducted Emissions (International) • EN55022-1998 - Radiated & Conducted Emissions (European Union) • EN55024 - 1998 - (Immunity) (European Union) • CE - EMC Directive (89/336/EEC) (European Union) • VCCI - Radiated & Conducted Emissions (Japan) • CNS13438 - Radiated & Conducted Emissions (Taiwan) • AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) • MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)

Intel® 10 Gigabit Network Adapter Specifications

Feature	Intel® Ethernet 10G 2P X540-t Adapter	Intel® Ethernet 10G 2P X520 Adapter	Intel® Ethernet Server Adapter X520-T2
Bus Connector	PCI Express 2.0	PCI Express 2.0	PCI Express 2.0
Bus Speed	x8	x8	x8
Transmission Mode/Connector	10GBase-T/RJ-45	Twinaxial copper/SFP+	10GBase-T/RJ-45
Cabling	10GBase-T (Category 6A)	10 Gigabit Ethernet over SFP+ Direct Attach Copper (10GSFP+Cu)	10GBase-T (Category 6A)
Power Requirements	15 W Maximum @ +12 V	6.2 W Maximum @ +3.3 V	25 W Maximum @ +12 V
Dimensions (excluding bracket)	5.7 x 2.7 in 14.5 x 6.9 cm	5.7 x 2.7 in 14.5 x 6.9 cm	6.59 x 2.71 in 16.7 x 6.9 cm

Feature	Intel® Ethernet 10G 2P X540-t Adapter	Intel® Ethernet 10G 2P X520 Adapter	Intel® Ethernet Server Adapter X520-T2
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)
MTBF	108 years	83.9 years	83.15 years
Available Speeds	10 Gbps/1 Gbps	10 Gbps/1 Gbps	10 Gbps/1 Gbps
Duplex Modes	Full only	Full only	Full only
Indicator Lights	*Two per port:* Link and Activity	*Two per port:* Link and Activity	Link Activity
Standards Conformance	IEEE 802.1p IEEE 802.1Q IEEE 802.3an IEEE 802.3ac IEEE 802.3ad IEEE 802.3an IEEE 802.3x ACPI v1.0 PCI Express 2.0	IEEE 802.1p IEEE 802.1Q IEEE 802.3an IEEE 802.3ac IEEE 802.3ad IEEE 802.3x ACPI v1.0 PCI Express 2.0	IEEE 802.1p IEEE 802.1Q IEEE 802.1ae IEEE 802.3ac IEEE 802.3ad IEEE 802.3an IEEE 802.3x ACPI v1.0 PCI Express 2.0
Regulatory and Safety	**Safety Compliance** • UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) • EN 60 950 (European Union) • IEC 60 950 (International) **EMC Compliance** • FCC Part 15 - Radiated & Conducted Emissions (USA) • ICES-003 - Radiated & Conducted Emissions (Canada) • CISPR 22 - Radiated & Conducted Emissions (International) • EN55022-1998 - Radiated & Conducted Emissions (European Union) • EN55024 - 1998 - (Immunity) (European Union) • CE - EMC Directive (89/336/EEC) (European Union) • VCCI - Radiated & Conducted Emissions (Japan) • CNS13438 - Radiated & Conducted Emissions (Taiwan) • AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) • MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)		

NOTE: For the Intel® 10 Gigabit AT Server Adapter, to ensure compliance with CISPR 24 and the EU's EN55024, this product should be used only with Category 6a shielded cables that are properly terminated according to the recommendations in EN50174-2.

Feature	Intel® Ethernet Server Adapter X520-2	Intel® Converged Network Adapter X710	Intel® Ethernet 10G 2P X550-t Adapter
Bus Connector	PCI Express 2.0	PCI Express 3.0	PCI Express 3.0
Bus Speed	x8	x8	x8
Transmission Mode/Connector	10GBase-SR/SFP+	SFP+	10GBase-T/RJ-45
Cabling	Multimode Fiber	Twinax 10GBase-SR/LR	10GBase-T (Category 6A)
Power Requirements	10.7 Watts @ +12 V	TBD	13W Maximum @ +12 V
Dimensions (excluding bracket)	5.73 x 2.71 in 14.6 x 6.9 cm	6.578 x 4.372 in 16.708 x 11.107 cm	5.13 x 2.7 in 13.0 x 6.9 cm

Feature	Intel® Ethernet Server Adapter X520-2	Intel® Converged Network Adapter X710	Intel® Ethernet 10G 2P X550-t Adapter
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)
MTBF	83.9 years	TBD	TBD
Available Speeds	10 Gbps/1 Gbps	10 Gbps/1 Gbps	10 Gbps/1 Gbps
Duplex Modes	Full only	Full Only	Full only
Indicator Lights	Link/Activity 1Gig/10Gig	Link/Activity 1Gig/10Gig	Link Activity
Standards Conformance	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3x ACPI v1.0 PCI Express 2.0	PCI Express 3.0 SFF-8431 IEEE 802.3z IEEE 802.3ae	IEEE 802.1p IEEE 802.1Q IEEE 802.3an IEEE 802.3ac IEEE 802.3ad IEEE 802.3x ACPI v1.0 PCI Express 3.0
Regulatory and Safety	**Safety Compliance** • UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) • EN 60 950 (European Union) • IEC 60 950 (International) **EMC Compliance** • FCC Part 15 - Radiated & Conducted Emissions (USA) • ICES-003 - Radiated & Conducted Emissions (Canada) • CISPR 22 - Radiated & Conducted Emissions (International) • EN55022-1998 - Radiated & Conducted Emissions (European Union) • EN55024 - 1998 - (Immunity) (European Union) • CE - EMC Directive (89/336/EEC) (European Union) • VCCI - Radiated & Conducted Emissions (Japan) • CNS13438 - Radiated & Conducted Emissions (Taiwan) • AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) • MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)		

Intel® 10 Gigabit Network Mezzanine Card Specifications

Feature	Intel® Ethernet X520 10GbE Dual Port KX4-KR Mezz	Intel® Ethernet X520 10GbE Dual Port KX4 Mezz
Bus Connector	PCI Express 2.0	PCI Express 2.0
Bus Speed	x8	x8
Power Requirements	7.4 Watts (maximum) @ 3.3 V	7.4 Watts (maximum) @ 3.3 V
Dimensions	3.65 x 3.3 in.	3.65 x 3.3 in.

Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)
MTBF	147 years	147 years
Available Speeds	10 Gbps/1 Gbps	10 Gbps/1 Gbps
Duplex Modes	Full only	Full only
Standards Conformance	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ae IEEE 802.3x ACPI v1.0 PCI Express 2.0	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ae IEEE 802.3x ACPI v1.0 PCI Express 2.0
Regulatory and Safety	**Safety Compliance** • UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) • EN 60 950 (European Union) • IEC 60 950 (International) **EMC Compliance** • FCC Part 15 - Radiated & Conducted Emissions (USA) • ICES-003 - Radiated & Conducted Emissions (Canada) • CISPR 22 - Radiated & Conducted Emissions (International) • EN55022-1998 - Radiated & Conducted Emissions (European Union) • EN55024 - 1998 - (Immunity) (European Union) • CE - EMC Directive (89/336/EEC) (European Union) • VCCI - Radiated & Conducted Emissions (Japan) • CNS13438 - Radiated & Conducted Emissions (Taiwan) • AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) • MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)	

Intel® 10GbE Network Daughter Cards (NDC) Specifications

Feature	Intel® Ethernet 10G 4P X540/I350 rNDC	Intel® Ethernet 10G 4P X520/I350 rNDC	Intel® Ethernet 10G 2P X520-k bNDC
Bus Connector	PCI Express 2.0	PCI Express 2.0	PCI Express 2.0
Bus Speed	x8	2 x8	x8
Transmission Mode/Connector	Twisted copper/RJ-45	SFP+	Copper/Backplane
Cabling	1000Base-T (Category 5 or Category 3 on 10 Mbps only)	SFP+ SR/DA	10GBase-KR and 1000Base-KX
Power Require-ments	5.5 Watts (maximum) @ 3.3 V	10.1 Watts (maximum) @ 12 V	0.6 Watts @ 3.3 V (AUX), 6.3 Watts @ 1.2 V (VCORE)

Dimensions	3.93 x 3.67 in.	4.3 x 3.7 in.	3.0 x 2.5 in.
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)
MTBF	68 years	65 years	147 years
Available Speeds	2 ports at 10 Gbps/(2 ports at 1 Gbps, see Intel® Ethernet Gigabit 4P X540/I350 rNDC)	2 ports at 10 Gbps/(2 ports at 1 Gbps, see Intel® Ethernet Gigabit 4P X520/I350 rNDC)	10 Gbps/1 Gbps
Duplex Modes	Full only	Full only	Full only
Standards Conformance	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ae IEEE 802.3x ACPI v1.0 PCI Express 1.0a	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ae IEEE 802.3x ACPI v1.0 PCI Express 1.0a	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ap IEEE 802.3x ACPI v1.0 PCI Express 2.0
Regulatory and Safety	**Safety Compliance** ● UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) ● EN 60 950 (European Union) ● IEC 60 950 (International) **EMC Compliance** ● FCC Part 15 - Radiated & Conducted Emissions (USA) ● ICES-003 - Radiated & Conducted Emissions (Canada) ● CISPR 22 - Radiated & Conducted Emissions (International) ● EN55022-1998 - Radiated & Conducted Emissions (European Union) ● EN55024 - 1998 - (Immunity) (European Union) ● CE - EMC Directive (89/336/EEC) (European Union) ● VCCI - Radiated & Conducted Emissions (Japan) ● CNS13438 - Radiated & Conducted Emissions (Taiwan) ● AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) ● MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)		

Feature	Intel® Ethernet 10G 4P X540/I350 rNDC	Intel® Ethernet 10G 4P X520/I350 rNDC	Intel® Ethernet 10G 2P X520-k bNDC
Bus Connector	PCI Express 2.0	PCI Express 2.0	PCI Express 2.0
Bus Speed	x8	2 x8	x8
Transmission Mode/Connector	Twisted copper/RJ-45	SFP+	Copper/Backplane
Cabling	1000Base-T (Category 5 or Category 3 on 10 Mbps only)	SFP+ SR/DA	10GBase-KR and 1000Base-KX

Power Requirements	5.5 Watts (maximum) @ 3.3 V	10.1 Watts (maximum) @ 12 V	0.6 Watts @ 3.3 V (AUX), 6.3 Watts @ 1.2 V (VCORE)
Dimensions	3.93 x 3.67 in.	4.3 x 3.7 in.	3.0 x 2.5 in.
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)
MTBF	68 years	65 years	147 years
Available Speeds	2 ports at 10 Gbps/(2 ports at 1 Gbps, see Intel® Ethernet Gigabit 4P X540/I350 rNDC)	2 ports at 10 Gbps/(2 ports at 1 Gbps, see Intel® Ethernet Gigabit 4P X520/I350 rNDC)	10 Gbps/1 Gbps
Duplex Modes	Full only	Full only	Full only
Standards Conformance	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ae IEEE 802.3x ACPI v1.0 PCI Express 1.0a	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ae IEEE 802.3x ACPI v1.0 PCI Express 1.0a	IEEE 802.1p IEEE 802.1Q IEEE 802.3ac IEEE 802.3ad IEEE 802.3ap IEEE 802.3x ACPI v1.0 PCI Express 2.0
Regulatory and Safety	**Safety Compliance** ● UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) ● EN 60 950 (European Union) ● IEC 60 950 (International) **EMC Compliance** ● FCC Part 15 - Radiated & Conducted Emissions (USA) ● ICES-003 - Radiated & Conducted Emissions (Canada) ● CISPR 22 - Radiated & Conducted Emissions (International) ● EN55022-1998 - Radiated & Conducted Emissions (European Union) ● EN55024 - 1998 - (Immunity) (European Union) ● CE - EMC Directive (89/336/EEC) (European Union) ● VCCI - Radiated & Conducted Emissions (Japan) ● CNS13438 - Radiated & Conducted Emissions (Taiwan) ● AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) ● MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)		

Feature	Intel® Ethernet 10G 4P X710-k bNDC	Intel® Ethernet 10G 4P X710/I350 rNDC	Intel® Ethernet 10G 4P X710 SFP+ rNDC
Bus Connector	Dell bNDC 13G	Dell bNDC 13G	Dell bNDC 13G
Bus Speed	x8	x8	x8
Transmission Mode/Connector	KX/KR	SFP+	SFP+
Cabling	Backplane	Cat-5e	Twinax 10GBase-SR/LR
Power Requirements	TBD	TBD	TBD
Dimensions	3.000x2.449 in 7.62x6.220cm	4.331x3.661 in 11.0x9.298 cm	4.331x3.661 in 11.0x9.298 cm
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)
MTBF	TBD	TBD	TBD
Available Speeds	1G/10G	1G/10G	1G/10G
Duplex Modes	Full only	Full only	Full only
Indicator Lights	None	Link/Activity Speed	Link/Activity Speed
Standards Conformance	PCI Express 3.0 IEEE 802.3ap	PCI Express 3.0 SFF-8431 IEEE 802.3z IEEE 802.3ae	PCI Express 3.0 SFF-8431 IEEE 802.3z IEEE 802.3ae
Regulatory and Safety	**Safety Compliance** • UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) • EN 60 950 (European Union) • IEC 60 950 (International) **EMC Compliance** • FCC Part 15 - Radiated & Conducted Emissions (USA) • ICES-003 - Radiated & Conducted Emissions (Canada) • CISPR 22 - Radiated & Conducted Emissions (International) • EN55022-1998 - Radiated & Conducted Emissions (European Union) • EN55024 - 1998 - (Immunity) (European Union) • CE - EMC Directive (89/336/EEC) (European Union) • VCCI - Radiated & Conducted Emissions (Japan) • CNS13438 - Radiated & Conducted Emissions (Taiwan) • AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) • MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)		

Intel® Gigabit Network Adapter Specifications

Feature	Intel® Gigabit 2P I350-t Adapter and Intel® Gigabit 4P I350-t Adapter
Bus Connector	PCI Express 2.0
Bus Speed	x4
Transmission Mode/Connector	Twisted copper/RJ-45
Cabling	1000Base-T (Category 3 or Category 5)
Power Requirements	*Intel® Gigabit 2P I350-t Adapter*: 4.8 Watts @ 12 V

	Intel® Gigabit 4P I350-t Adapter: 6.0 Watts @ 12 V
Dimensions (excluding bracket)	5.3 x 2.7 in. 13.5 x 6.9 cm
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)

MTBF	68 years
Available Speeds	10/100/1000 auto-negotiate
Duplex Modes	Full or half at 10/100 Mbps; full only at 1000 Mbps
Standards Conformance	IEEE 802.1p IEEE 802.1Q IEEE 802.3ab IEEE 802.3ac IEEE 802.3ad IEEE 802.3az IEEE 802.3u IEEE 802.3x IEEE 802.3z ACPI v1.0 PCI Express 2.0
Indicator Lights	*Two per port:* Activity and Speed
Regulatory and Safety	**Safety Compliance** ● UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) ● EN 60 950 (European Union) ● IEC 60 950 (International) **EMC Compliance** ● FCC Part 15 - Radiated & Conducted Emissions (USA) ● ICES-003 - Radiated & Conducted Emissions (Canada) ● CISPR 22 - Radiated & Conducted Emissions (International) ● EN55022-1998 - Radiated & Conducted Emissions (European Union) ● EN55024 - 1998 - (Immunity) (European Union) ● CE - EMC Directive (89/336/EEC) (European Union) ● VCCI - Radiated & Conducted Emissions (Japan) ● CNS13438 - Radiated & Conducted Emissions (Taiwan) ● AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) ● MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)

Intel® Gigabit Network Mezzanine Card Specifications

Feature	Intel® Gigabit 4P I350-t Mezz
Bus Connector	PCI Express 2.0
Bus Speed	x4
Power Requirements	3.425 Watts (maximum) @ 3.3 V
Dimensions	3.65 x 3.3 in.
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)
MTBF	108 years
Available Speeds	Full only at 1000 Mbps
Duplex Modes	Full at 1000 Mbps
Standards Conformance	IEEE 802.1p IEEE 802.1Q IEEE 802.3ab IEEE 802.3ac IEEE 802.3ad IEEE 802.3x ACPI v1.0 PCI Express 2.0

Regulatory and Safety	Safety Compliance ● UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) ● EN 60 950 (European Union) ● IEC 60 950 (International) **EMC Compliance** ● FCC Part 15 - Radiated & Conducted Emissions (USA) ● ICES-003 - Radiated & Conducted Emissions (Canada) ● CISPR 22 - Radiated & Conducted Emissions (International) ● EN55022-1998 - Radiated & Conducted Emissions (European Union) ● EN55024 - 1998 - (Immunity) (European Union) ● CE - EMC Directive (89/336/EEC) (European Union) ● VCCI - Radiated & Conducted Emissions (Japan) ● CNS13438 - Radiated & Conducted Emissions (Taiwan) ● AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) ● MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)

Intel® Gigabit Network Daughter Cards Specifications

Feature	Intel® Gigabit 4P X710/I350 rNDC	Intel® Gigabit 4P I350 bNDC
Bus Connector	Dell bNDC 13G	Dell rNDC 13G
Bus Speed	x2	x4
Transmission Mode/Connector	1000Base-T	KX
Cabling	Backplane	Backplane
Power Requirements	TBD	TBD
Dimensions (excluding bracket)	4.331 x 3.661 in 11.007 x 9.298 cm	3.000 x 2.449 in 7.620 x 6.220 cm
Operating Temperature	32 - 131 deg. F (0 - 55 deg. C)	32 - 131 deg. F (0 - 55 deg. C)
MTBF	TBD	TBD
Available Speeds	10/100/1000	1G
Duplex Modes	Full	Full
Indicator Lights	*Two per port:* Link/Activity Speed	None
Standards Conformance	PCI Express 2.1 IEEE 802.3i IEEE 802.3ab IEEE 802.3u IEEE 802.3ad IEEE 802.3az	PCI Express 3.0 IEEE 802.3ap
Regulatory and Safety	**Safety Compliance** ● UL 60950 Third Edition- CAN/CSA-C22.2 No.60950-00 (USA/Canada) ● EN 60 950 (European Union) ● IEC 60 950 (International) **EMC Compliance** ● FCC Part 15 - Radiated & Conducted Emissions (USA) ● ICES-003 - Radiated & Conducted Emissions (Canada)	

	• CISPR 22 - Radiated & Conducted Emissions (International) • EN55022-1998 - Radiated & Conducted Emissions (European Union) • EN55024 - 1998 - (Immunity) (European Union) • CE - EMC Directive (89/336/EEC) (European Union) • VCCI - Radiated & Conducted Emissions (Japan) • CNS13438 - Radiated & Conducted Emissions (Taiwan) • AS/NZS3548 - Radiated & Conducted Emissions (Australia/New Zealand) • MIC notice 1997-41, EMI and MIC notice 1997-42 - EMS (Korea)

Standards

- IEEE 802.1p: Priority Queuing (traffic prioritizing) and Quality of Service levels
- IEEE 802.1Q: Virtual LAN identification
- IEEE 802.3ab: Gigabit Ethernet over copper
- IEEE 802.3ac: Tagging
- IEEE 802.3ad: SLA (FEC/GEC/Link Aggregation - static mode)
- IEEE 802.3ad: Dynamic mode
- IEEE 802.3ae: 10 Gbps Ethernet
- IEEE 802.3an: 10GBase-T 10 Gbps Ethernet over unshielded twisted pair
- IEEE 802.3ap: Backplane Ethernet
- IEEE 802.3u: Fast Ethernet
- IEEE 802.3x: Flow Control
- IEEE 802.3z: Gigabit Ethernet over optical fiber
- ACPI: Advanced Configuration and Power Management
- PCI Express: system bus specification: 32/64-bit, x1, x2, x4, x8, x16

More information on IEEE 802 standards available at http://www.ieee802.org.

IEEE 802.3ac VLANs:

VLANs require VLAN-capable switches either implicit (switch only) or explicit (IEEE 802.3ac). IEEE 802.3ac VLANs allow multiple VLANs per adapter or team since both the switch and adapter use a tag in the packet header to sort VLANs.

Intel gigabit and 10 gigabit network adapters fully support implicit and explicit VLANs.

Software License Agreement

INTEL SOFTWARE LICENSE AGREEMENT (Final, License)

IMPORTANT - READ BEFORE COPYING, INSTALLING OR USING.

Do not use or load this software and any associated materials (collectively, the "Software") until you have carefully read the following terms and conditions. By loading or using the Software, you agree to the terms of this Agreement. If you do not wish to so agree, do not install or use the Software.

LICENSES

Please Note:

- If you are a network administrator, the "Site License" below shall apply to you.
- If you are an end user, the "Single User License" shall apply to you.

SITE LICENSE. You may copy the Software onto your organization's computers for your organization's use, and you may make a reasonable number of back-up copies of the Software, subject to these conditions:

1. **This Software is licensed for use only in conjunction with Intel component products. Use of the Software in conjunction with non-Intel component products is not licensed hereunder.**
2. You may not copy, modify, rent, sell, distribute or transfer any part of the Software except as provided in this Agreement, and you agree to prevent unauthorized copying of the Software.
3. You may not reverse engineer, decompile, or disassemble the Software.
4. You may not sublicense or permit simultaneous use of the Software by more than one user.
5. The Software may include portions offered on terms in addition to those set out here, as set out in a license accompanying those portions.

SINGLE USER LICENSE. You may copy the Software onto a single computer for your personal, noncommercial use, and you may make one back-up copy of the Software, subject to these conditions:

1. **This Software is licensed for use only in conjunction with Intel component products. Use of the Software in conjunction with non-Intel component products is not licensed hereunder.**
2. You may not copy, modify, rent, sell, distribute or transfer any part of the Software except as provided in this Agreement, and you agree to prevent unauthorized copying of the Software.
3. You may not reverse engineer, decompile, or disassemble the Software.
4. You may not sublicense or permit simultaneous use of the Software by more than one user.
5. The Software may include portions offered on terms in addition to those set out here, as set out in a license accompanying those portions.

OWNERSHIP OF SOFTWARE AND COPYRIGHTS. Title to all copies of the Software remains with Intel or its suppliers. The Software is copyrighted and protected by the laws of the United States and other countries, and international treaty provisions. You may not remove any copyright notices from the Software. Intel may make changes to the Software, or to items referenced therein, at any time without notice, but is not obligated to support or update the Software. Except as otherwise expressly provided, Intel grants no express or implied right under Intel patents, copyrights, trademarks, or other intellectual property rights. You may transfer the Software only if the recipient agrees to be fully bound by these terms and if you retain no copies of the Software.

LIMITED MEDIA WARRANTY. If the Software has been delivered by Intel on physical media, Intel warrants the media to be free from material physical defects for a period of ninety days after delivery by Intel. If such a defect is found, return the media to Intel for replacement or alternate delivery of the Software as Intel may select.

EXCLUSION OF OTHER WARRANTIES. EXCEPT AS PROVIDED ABOVE, THE SOFTWARE IS PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A PARTICULAR PURPOSE. Intel does not warrant or assume responsibility for the accuracy or completeness of any information, text, graphics, links or other items contained within the Software.

LIMITATION OF LIABILITY. IN NO EVENT SHALL INTEL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION, OR LOST INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF INTEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SOME JURISDICTIONS PROHIBIT EXCLUSION OR LIMITATION OF LIABILITY FOR IMPLIED WARRANTIES OR CONSEQUENTIAL OR INCIDENTAL DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU. YOU MAY ALSO HAVE OTHER LEGAL RIGHTS THAT VARY FROM JURISDICTION TO JURISDICTION.

TERMINATION OF THIS AGREEMENT. Intel may terminate this Agreement at any time if you violate its terms. Upon termination, you will immediately destroy the Software or return all copies of the Software to Intel.

APPLICABLE LAWS. Claims arising under this Agreement shall be governed by the laws of California, excluding its principles of conflict of laws and the United Nations Convention on Contracts for the Sale of Goods. You may not export the Software in violation of applicable export laws and regulations. Intel is not obligated under any other agreements unless they are in writing and signed by an authorized representative of Intel.

GOVERNMENT RESTRICTED RIGHTS. The Software is provided with "RESTRICTED RIGHTS." Use, duplication, or disclosure by the Government is subject to restrictions as set forth in FAR52.227-14 and DFAR252.227-7013 *et seq*. or its successor. Use of the Software by the Government constitutes acknowledgment of Intel's proprietary rights therein. Contractor or Manufacturer is Intel.