Intel® Ethernet Adapters and Devices User Guide

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.

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 5 .

NOTE: If you update the firmware, you must update the driver software to the same family version.

1. Make sure that you are installing the latest driver software for your adapter. Visit Intel's support website to download the latest drivers.
2. Review system requirements.
3. Insert the adapter(s) in the computer.
4. Attach the copper or fiber network cable(s).
5. Install the driver.
6. For Windows systems, install the Intel® PROSet software.

If you have any problems with basic installation, see Troubleshooting.

You can now set up advanced features, if necessary. The available features and the configuration process varies with the adapter and your operating system.

Intended Audience

This document is intended for information technology professionals with a high level of knowledge, experience, and competency in Ethernet networking technology.

Related Documentation

Document	Description	Location
readme.txt	Release-specific information, important notes, and known issues.	Root folder of software download package
license.pdf	Licenses pertaining to this software release.	Root folder of software download package
Autorun.exe	Executable that launches the software package installer, which also provides links to the user guide and release notes.	Root folder of software download package

Before You Begin

Supported Devices

For help identifying your network device and finding supported devices, click the link below:

http://www.intel.com/support

Compatibility Notes

In order for an adapter based on the XL710 controller to reach its full potential, you must install it in a PCIe Gen3 x8 slot. Installing it in a shorter slot, or a Gen2 or Gen1 slot, will limit the throughput of the adapter.

Some older Intel® Ethernet Adapters do not have full software support for the most recent versions of Microsoft Windows*. Many older Intel Ethernet Adapters have base drivers supplied by Microsoft Windows. Lists of supported devices per OS are available at http://www.intel.com/support/go/network/adapter/nicoscomp.htm

Supported Operating Systems

Supported 32-bit Operating Systems

NOTE: Microsoft* Windows* 32-bit operating systems are only supported on Intel 1GbE Ethernet Adapters and slower devices. All adapters support 32-bit versions of Linux* and FreeBSD*.

Basic software and drivers are supported on the following operating systems:

- DOS
- SunSoft* Solaris* (drivers and support are provided by the operating system vendor)

Advanced software and drivers are supported on the following operating systems:

- Microsoft Windows 7
- Microsoft Windows 8.1
- Microsoft Windows 10
- Microsoft Windows 10 Version 1809
- Linux*, v2.4 kernel or higher
- FreeBSD*

Supported Intel® 64 Architecture Operating Systems

- Microsoft* Windows* 7
- Microsoft Windows 8.1
- Microsoft Windows 10
- Microsoft* Windows Server* 2008 R2
- Microsoft Windows Server 2012
- Microsoft Windows Server 2012 R2
- Microsoft Windows Server 2016
- Microsoft Windows Server 2016 Nano Server
- Microsoft Windows Server 2019

- VMWare ESXi 5.5
- VMWare* ESXi* 6.0
- VMWare* ESXi* 6.5 U1‡
- Ubuntu 14.04
- Red Hat* Linux*
- Novell* SUSE* Linux
- FreeBSD*

‡ If you upgrade to VMWare ESXi 6.5 U1 or later, you must use native mode drivers (available in the latest driver software package) instead of VMKLinux legacy mode drivers.

Supported Operating Systems for Itanium-based Systems

- Linux, v2.x kernel and higher, except v2.6

Some older Intel® Ethernet Adapters do not have full software support for the most recent versions of Microsoft Windows*. Many older Intel Ethernet Adapters have base drivers supplied by Microsoft Windows. Lists of supported devices per OS are available at http://www.intel.com/support/go/network/adapter/nicoscomp.htm

Hardware Compatibility

Before installing the adapter, check your system for the following:

- The latest BIOS for your system
- One open PCI Express slot

NOTE: The Intel® 10 Gigabit AT Server Adapter will only fit into x8 or larger PCI Express slots. Some systems have physical x8 PCI Express slots that actually support lower speeds. Please check your system manual to identify the slot.

Cabling Requirements

RJ-45 Connector Cables

The following table shows the maximum lengths for each cable type at a given transmission speed.

	Category 5	Category 6	Category 6a	Category 7
1 Gbps	100m	100m	100m	100m
10 Gbps	NA	55m	100m	100m
25 Gbps	NA	NA	NA	50m
40 Gbps	NA	NA	NA	50m

NOTE: To insure compliance with CISPR 24 and the EU's EN55024, devices based on the 82576 controller should be used only with CAT 5E shielded cables that are properly terminated according to the recommendations in EN50174-2.

Fiber Optic Cables (LC and SC connector types)

Intel Ethernet LR Server Adapters

Laser wavelength: 1310 nanometer (not visible)

Connector type: LC

Cable type: Single-mode fiber with 9.0µm core diameter

- Maximum cable length: 10 kilometers

Intel Ethernet 10 Gigabit SR Server Adapters

Laser wavelength: 850 nanometer (not visible)

Connector type: LC or SC

Cable type: Multi-mode fiber with 62.5µm core diameter

- Maximum cable length: 33 meters

Cable type: Multi-mode fiber with 50µm core diameter

- Maximum cable length: 300 meters

Intel Ethernet Gigabit SR Server Adapters

Laser wavelength: 850 nanometer (not visible)

Connector type: LC or SC

Cable type: Multi-mode fiber with 62.5µm core diameter

- Maximum cable length: 275 meters

Cable type: Multi-mode fiber with 50µm core diameter

- Maximum cable length: 550 meters

SFP+ Module and Direct Attach Cables

Consult your module specifications.

Installation Overview

Installing the 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 slot and secure the bracket to the chassis.
4. Replace the computer cover, then plug in the power cord.

Install Drivers and Software

Windows* Operating Systems

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

1. Download the latest drivers from the support website and transfer them to the system.
2. If the Found New Hardware Wizard screen is displayed, click **Cancel**.
3. Start the autorun located in the downloaded the software package. The autorun may automatically start after you have extracted the files.

4. Click **Install Drivers and Software**.
5. Follow the instructions in the install wizard.

Installing Linux* Drivers from Source Code

1. Download and expand the base driver tar file.
2. Compile the driver module.
3. Install the module using the modprobe command.
4. Assign an IP address using the ifconfig command.

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 network driver configuration, operating system or system BIOS. 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.

General Optimization

- Install the adapter in an appropriate slot.

  NOTE: Some PCIe x8 slots are actually configured as x4 slots. These slots have insufficient bandwidth for full line rate with some dual port 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.

- In order for an Intel® 710 Series Network Adapter to reach its full potential, you must install it in a PCIe Gen3 x8 slot. Installing it in a shorter slot, or a Gen2 or Gen1 slot, will impact the throughput the adapter can attain.
- Use the proper cabling for your device.
- Enable Jumbo Packets, if your other network components can also be configured for it.

- 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.
- 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.
- When passing traffic on multiple network ports using an I/O application that runs on most or all of the cores in your system, consider setting the CPU Affinity for that application to fewer cores. This should reduce CPU utilization and in some cases may increase throughput for the device. The cores selected for CPU Affinity must be local to the affected network device's Processor Node/Group. You can use the PowerShell command Get-NetAdapterRSS to list the cores that are local to a device. You may need to increase the number of cores assigned to the application to maximize throughput. Refer to your operating system documentation for more details on setting the CPU Affinity.
- If you have multiple 10 Gpbs (or faster) ports installed in a system, the RSS queues of each adapter port can be adjusted to use non-overlapping sets of processors within the adapter's local NUMA Node/Socket. Change the RSS Base Processor Number for each adapter port so that the combination of the base processor and the max number of RSS processors settings ensure non-overlapping cores.
    1. Identify the adapter ports to be adjusted and inspect at their RssProcessorArray using the Get-NetAdapterRSS PowerShell cmdlet.
    2. Identify the processors with NUMA distance 0. These are the cores in the adapter's local NUMA Node/Socket and will provide the best performance.
    3. Adjust the RSS Base processor on each port to use a non-overlapping set of processors within the local set of processors. You can do this manually or using the following PowerShell command:
       ```
       Set-NetAdapterAdvancedProperty -Name <Adapter Name> -DisplayName
       "RSS Base
       Processor Number" -DisplayValue <RSS Base Proc Value>
       ```
    4. Use the Get-NetAdpaterAdvancedproperty cmdlet to check that the right values have been set:
       ```
       Get-NetAdpaterAdvancedproperty -Name <Adapter Name>
       ```

       For Example: For a 4 port adapter with Local processors 0, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, and 'Max RSS processor' of 8, set the RSS base processors to 0, 8, 16 and 24.

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.

Remote Storage

The remote storage features allow you to access a SAN or other networked storage using Ethernet protocols. This includes Data Center Bridging (DCB), iSCSI over DCB, and Fibre Channel over Ethernet (FCoE).

DCB (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.



iSCSI Over DCB

Intel® Ethernet adapters support iSCSI software initiators that are native to the underlying operating system. In the case of Windows, the Microsoft iSCSI Software Initiator, enables connection of a Windows host to an external iSCSI storage array using an Intel Ethernet adapter.

In the case of Open Source distributions, virtually all distributions include support for an Open iSCSI Software Initiator and Intel® Ethernet adapters will support them. Please consult your distribution documentation for additional configuration details on their particular Open iSCSI initiator.

Intel® 82599 and X540-based adapters support iSCSI within a Data Center Bridging cloud. Used in conjunction with switches and targets that support the iSCSI/DCB application TLV, this solution can provide guaranteed minimum bandwidth for iSCSI traffic between the host and target. This solution enables storage administrators to segment iSCSI traffic from LAN traffic, similar to how they can currently segment FCoE from LAN traffic. Previously, iSCSI traffic within a DCB supported environment was treated as LAN traffic by switch vendors. Please consult your switch and target vendors to ensure that they support the iSCSI/DCB application TLV.

Intel® Ethernet FCoE (Fibre Channel over Ethernet)

Fibre Channel over Ethernet (FCoE) is the encapsulation of standard Fibre Channel (FC) protocol frames as data within standard Ethernet frames. This link-level encapsulation, teamed with an FCoE-aware Ethernet-to-FC gateway, acts to extend an FC fabric to include Ethernet-based host connectivity. The FCoE specification focuses on encapsulation of FC frames specific to storage class traffic, as defined by the Fibre Channel FC-4 FCP specification.

NOTES:

- FCoE is no longer supported and is no longer being tested. Installing drivers from Release 23.2 or later may cause FCoE to behave unexpectedly or cease functioning.
- Support for new operating systems will not be added to FCoE. The last operating system versions that support FCoE are as follows:
  - Microsoft* Windows Server* 2012 R2
  - RHEL 7.2
  - RHEL 6.7
  - SLES 12 SP1
  - SLES 11 SP4
  - VMware* ESX 6.0 U3

Jumbo Frames

The base driver supports FCoE mini-Jumbo Frames (2.5k bytes) independent of the LAN Jumbo Frames setting.

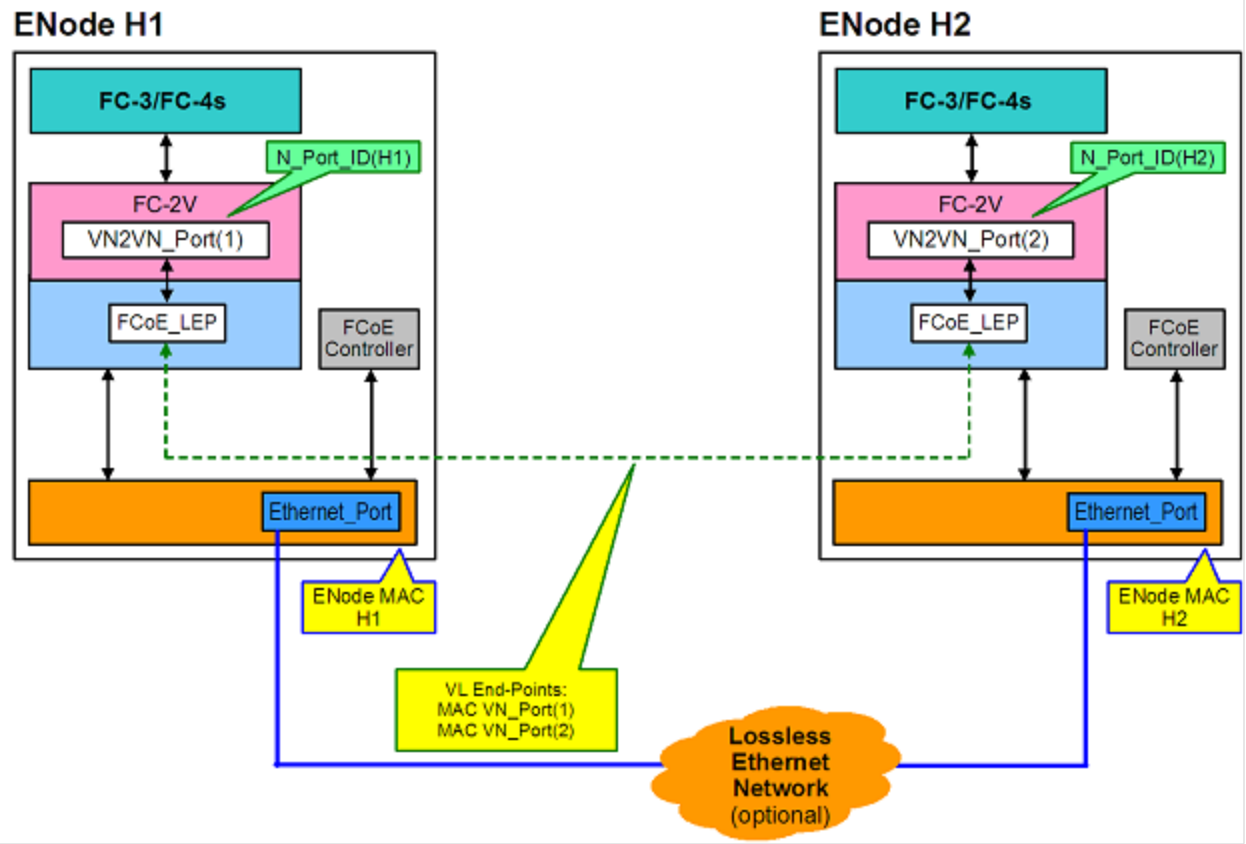
FCoE VN to VN (VN2VN) Support

FCoE VN to VN, also called VN2VN, is a standard for connecting two end-nodes (ENodes) directly using FCoE. An ENode can create a VN2VN virtual link with another remote ENode by not connecting to FC or FCoE switches (FCFs) in between, so neither port zoning nor advance fibre channel services is required. The storage software controls access to, and security of, LUNs using LUN masking. The VN2VN fabric may have a lossless Ethernet switch between the ENodes. This allows multiple ENodes to participate in creating more than one VN2VN virtual link in the VN2VN fabric. VN2VN has two operational modes: Point to Point (PT2PT) and Multipoint.

NOTE: The mode of operation is used only during initialization.

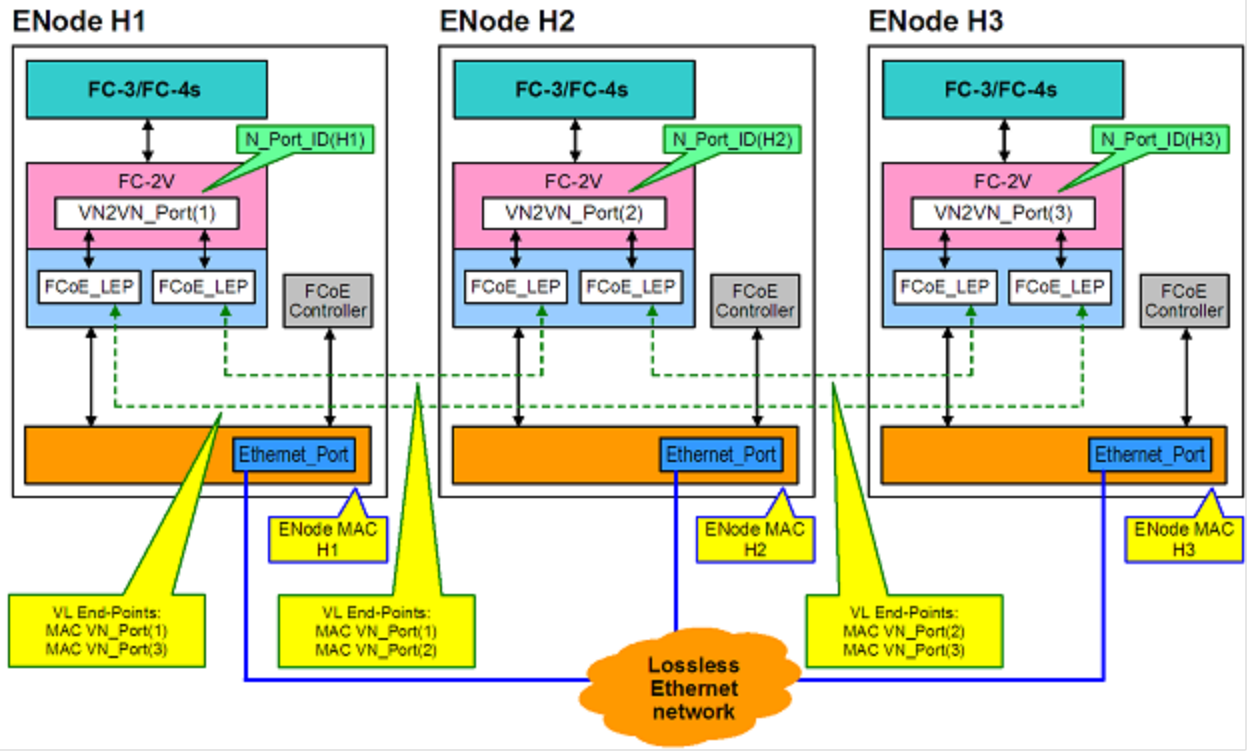
Point to Point (PT2PT) Mode

In Point to Point mode, there are only two ENodes, and they are connected either directly or through a lossless Ethernet switch:



MultiPoint Mode

If more than two ENodes are detected in the VN2VN fabric, then all nodes should operate in Multipoint mode:



Enabling VN2VN in Microsoft Windows

To enable VN2VN in Microsoft Windows:

1. Start Windows Device Manager.
2. Open the appropriate FCoE miniport property sheet (generally under Storage controllers) and click on the Advanced tab.
3. Select the VN2VN setting and choose "Enable."

Remote Boot

Remote Boot allows you to boot a system using only an Ethernet adapter. You connect to a server that contains an operating system image and use that to boot your local system.

Intel® Boot Agent

The Intel® Boot Agent is a software product that allows your networked client computer to boot using a program code image supplied by a remote server. Intel Boot Agent complies with the Pre-boot eXecution Environment (PXE) Version 2.1 Specification. It is compatible with legacy boot agent environments that use BOOTP protocol.

Supported Devices

Intel Boot Agent supports all Intel 10 Gigabit Ethernet, 1 Gigabit Ethernet, and PRO/100 Ethernet Adapters.

Intel® Ethernet iSCSI Boot

Intel® Ethernet iSCSI Boot provides the capability to boot a client system from a remote iSCSI disk volume located on an iSCSI-based Storage Area Network (SAN).

NOTE: Release 20.6 is the last release in which Intel® Ethernet iSCSI Boot supports Intel® Ethernet Desktop Adapters and Network Connections. Starting with Release 20.7, Intel Ethernet iSCSI Boot no longer supports Intel Ethernet Desktop Adapters and Network Connections.

Intel® Ethernet FCoE Boot

Intel® Ethernet FCoE Boot provides the capability to boot a client system from a remote disk volume located on an Fibre Channel Storage Area Network (SAN).

Using Intel® PROSet for Windows Device Manager

There are two ways to navigate to the FCoE properties in Windows Device Manager: by using the "Data Center" tab on the adapter property sheet or by using the Intel® "Ethernet Virtual Storage Miniport Driver for FCoE Storage Controllers" property sheet.

Supported Devices

A list of Intel Ethernet Adapters that support FCoE can be found at http://www.intel.com/support/go/network/adapter/fcoefaq.htm

Virtualization Support

Virtualization 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, 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.
- When you attach a Virtual Machine to a tenant overlay network through the Virtual NIC ports on a Virtual Switch, the encapsulation headers increase the Maximum Transmission Unit (MTU) size on the virtual port. The Encapsulation Overhead feature automatically adjusts the physical port's MTU size to compensate for this increase.
- 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 Microsoft* 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.
- While Microsoft supports Hyper-V on the Windows* 8.1 client OS, Intel® Ethernet adapters do not support virtualization settings (VMQ, SR-IOV) on Windows 8.1 client.
- 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® PROSet will automatically enable VMQ in the virtual machine switch. For driver-only installations, you must manually enable VMQ in the virtual machine switch.

NOTE: Intel ANS VLANs are not compatible with the Microsoft Hyper-V virtual machine switch. If you want to bind the virtual machine switch to a VLAN, you must create the VLAN from within the Virtual Switch Manager.

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.

NOTE: This step is not required for the team. When the Virtual NIC is created, its protocols are correctly bound.

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.

NOTE: Support for the Intel PROSet command line utilities (prosetcl.exe and crashdmp.exe) has been removed, and is no longer installed. This functionality has been replaced by the Intel Netcmdlets for Microsoft* Windows PowerShell*. Please transition all of your scripts and processes to use the Intel Netcmdlets for Microsoft Windows PowerShell.

The following is an example of how to set up the configuration using Microsoft* Windows PowerShell*.

1. Get all the adapters on the system and store them into a variable.

```
$a = Get-IntelNetAdapter
```

2. Create a team by referencing the indexes of the stored adapter array.

```
New-IntelNetTeam -TeamMembers $a[1],$a[2] -TeamMode
 VirtualMachineLoadBalancing -TeamName "Team1"
```

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

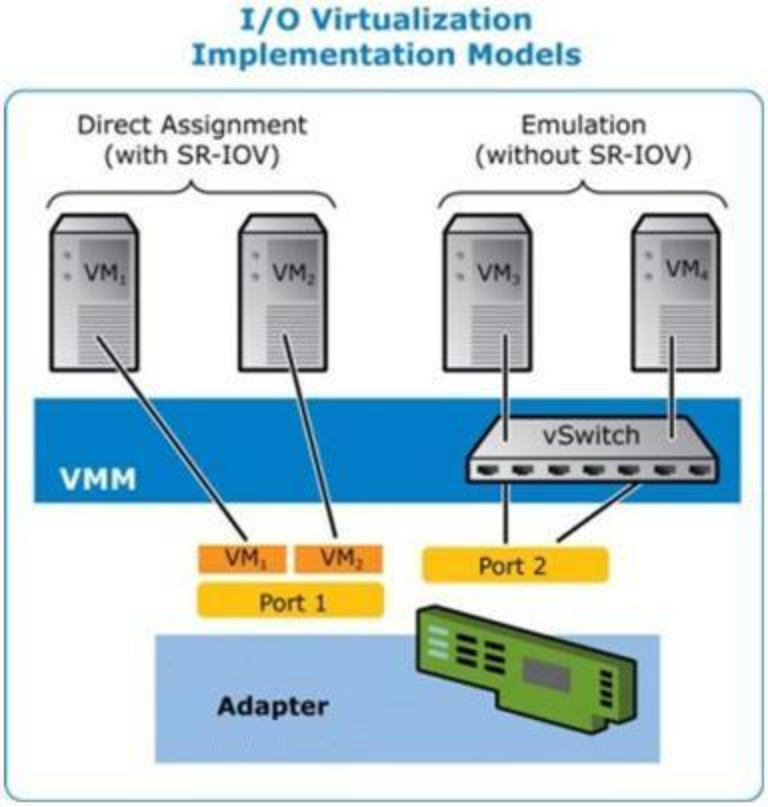
- If VMQ is not enabled for all adapters in a team, VMQ will be disabled for the team.
- If an adapter that does not support VMQ is added to a team, VMQ will be disabled for the team.
- Virtual NICs cannot be created on a team with Receive Load Balancing enabled. Receive Load Balancing is automatically 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 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.

NOTES:

- SR-IOV must be enabled in the BIOS.
- In Windows Server 2012, SR-IOV is not supported with teaming and VLANS. This occurs because the Hyper-V virtual switch does not enable SR-IOV on virtual interfaces such as teaming or VLANs. To enable SR-IOV, remove all teams and VLANs.

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

iWARP RDMA for Intel Ethernet

Remote Direct Memory Access, or RDMA, allows a computer to access another computer's memory without interacting with either computer's operating system data buffers, thus increasing networking speed and throughput. iWARP is a protocol for implementing RDMA across Internet Protocol networks.

Microsoft* Windows* provides two forms of RDMA: Network Direct Kernel (NDK) and Network Direct (ND).

iWARP RDMA for NDK Features such as SMB Direct (Server Message Block)

NDK allows Windows components (such as SMB Direct storage) to use iWARP features.

iWARP NDK Configuration

NDK functionality is included in the Intel base networking drivers and requires no additional features to be installed.

Verifying iWARP RDMA operation

Check RDMA is enabled on the network interfaces

Get-NetAdapterRDMA

Verify that the network interfaces show as RDMA capable and multichannel is enabled

Get-SmbClientNetworkInterface

Network Direct is enabled in the OS

Get-NetOffloadGlobalSetting | Select NetworkDirect

Each RDMA-capable network interface should have a listener at port 445 (Windows Client OSs that support RDMA may not post listeners)

netstat.exe -xan | ? {$_ -match "445"}

iWARP Routing Across IP Subnets

If you want to allow NDK's RDMA feature across subnets, you will need to select "Enable iWARP routing across IP Subnets" on the iWARP Configuration Options screen during base driver installation (see Installation below).

iWARP for Virtualized Environments

In order to enable RDMA capability on virtual adapter(s) connected to a VMSwitch, *SRIOV (Single Root IO Virtualization) and *VMQ (Virtual Machine Queues) advanced properties must be enabled on each port driver. Under certain circumstances, these settings may be disabled by default. These options can be set manually in the advanced tab of the adapter properties dialog box, or the following Powershell commands may be used:

```
Set-NetAdapterAdvancedProperty -Name <nic_name> -RegistryKeyword *SRIOV -
RegistryValue 1
```

```
Set-NetAdapterAdvancedProperty -Name <nic_name> -RegistryKeyword *VMQ -
RegistryValue 1
```

Testing iWARP NDK: Microsoft* Windows* SMB Direct with DiskSPD

This section outlines a recommended way to test iWARP RDMA for Intel Ethernet functionality and performance on Microsoft* Windows* operating systems.

Note that since SMB Direct is a storage workload, the performance of the benchmark may be limited to the speed of the storage device rather than the network interface being tested. Intel recommends using the fastest storage possible in order to test the true capabilities of the network device(s) under test.

Test instructions:

1. Set up at least two Microsoft* Windows* Server 2012 R2 or Microsoft* Windows* Server 2016 servers with at least one network interface of iWARP-capable Intel® Ethernet connected per server.
2. On the system designated as the SMB server, set up an SMB share. Note that the performance of the benchmark may be limited to the speed of the storage device rather than the network interface being tested. Storage setup is outside of the scope of this document.
   PowerShell (PS) command:
   New-SmbShare -Name <SMBsharename> -Path <SMBsharefilepath> -FullAccess <domain-name>\Administrator,Everyone

   Example:
   New-SmbShare -Name RAMDISKShare -Path R:\RAMDISK -FullAccess group\Administrator,Everyone
3. Download and install Microsoft utility Diskspd from here: https://gallery.technet.microsoft.com/DiskSpd-a-robust-storage-6cd2f223
4. Using CMD or Powershell, cd to the DiskSpd folder and run tests. (Refer to Diskspd documentation for more details on parameters)

   Example: Set the block size to 4K, run the test for 60 seconds, disable all hardware and software caching, measure and display latency statistics, leverage 16 overlapped IOs and 16 threads per target, random 0% writes and 100% reads and create a 10GB test file at "\\<SMBserverX722IP>\<SMBsharename>\test.dat" :
   .\diskspd.exe –b4K -d60 -h -L –o16 -t16 –r –w0 -c10G \\<SMBserver-X722IP>\<SMBsharename>\test.dat
5. Verify RDMA traffic is running using perfmon counters such as "RDMA Activity" and "SMB Direct Connection". Refer to Microsoft documentation for more details.

Configuring iWARP Guest Support (NDK Mode 3)

NDK Mode 3 allows kernel mode Windows components to use iWARP features inside Hyper-V guest partitions. To enable NDK mode 3 on an Intel Ethernet device, do the following:

1. Enable SR-IOV in your system's BIOS or uEFI.
2. Enable the SR-IOV advanced option on the device.
3. Configure the device with SR-IOV enabled on the switch. Perform this step for all functions on the same device:
   ```
   New-VMSwitch -Name <switch_name> -NetAdapterName <device_name>
   -EnableIov $true
   ```
4. Configure the number of RDMA virtual functions (VFs) on the device by setting the "RdmaMaxVfsEnabled" advanced option. All physical functions must be set to the same value. The value is the maximum number of VFs that can be capable of RDMA at one time for the entire device. Enabling more VFs will restrict RDMA resources from physicals functions (PFs) and other VFs.
   ```
   Set-NetAdapterAdvancedProperty -Name <device_name> -RegistryKeyword
   RdmaMaxVfsEnabled -RegistryValue <Value: 0 – 32>
   ```

5. Disable all PF adapters on the host and re-enable them. This is required when the registry keyword "RdmaMaxVfsEnabled" is changed or when creating or destroying a VMSwitch.
   ```
   Get-NetAdapterRdma | Disable-NetAdapter
   Get-NetAdapterRdma | Enable-NetAdapter
   ```
6. Create VM Network Adapters for VMs that require RDMA VF support.
   ```
   Add-VMNetworkAdapter -VMName <vm_name> -VMNetworkAdapterName <device_
   name> -SwitchName <switch_name>
   ```
7. If you plan to use Microsoft Windows 10 Creators Update (RS2) or later on a guest partition, set the RDMA weight on the VM Network Adapter by entering the following command on the host:
   ```
   Set-VMNetworkAdapterRdma -VMName <vm_name> -VMNetworkAdapterName <device_
   name> -RdmaWeight 100
   ```
8. Set SR-IOV weight on the VM Network Adapter (Note: SR-IOV weight must be set to 0 before setting the RdmaWeight to 0):
   ```
   Set-VMNetworkAdapter -VMName <vm_name> -VMNetworkAdapterName <device_
   name> -IovWeight 100
   ```
9. Install the VF network adapter with the PROSET Installer in the VM.
10. Enable RDMA on the VF driver and Hyper-V Network Adapter using PowerShell in the VM:
    ```
    Set-NetAdapterAdvancedProperty -Name <device_name> -RegistryKeyword
    RdmaVfEnabled -RegistryValue 1
    Get-NetAdapterRdma | Enable-NetAdapterRdma
    ```

iWARP RDMA for ND User-Mode Applications

ND allows user-mode applications to use iWARP features. If you plan to make use of iWARP features in user-mode applications you are developing, you will need to install the Network Direct (ND) feature when you install the drivers.

iWARP User Mode Installation

The Intel® Ethernet User Mode iWARP Provider is supported on Microsoft* Windows Server* 2012 R2 or later.

NOTES:

- For installation on Windows Server 2016 Nano Server, see Installing on Nano Server below.
- User mode applications may have prerequisites such as Microsoft HPC Pack or Intel MPI Library, refer to your application documentation for more details.

Follow the steps below to install user-mode Network Direct (ND) iWARP features.

1. From the installation media, run Autorun.exe to launch the installer, then choose "Install Drivers and Software" and accept the license agreement.
2. On the Setup Options screen, select "Intel® Ethernet User Mode iWARP Provider".
3. On the iWARP Configuration Options screen, select "Enable iWARP routing across IP Subnets" if desired. Note that this option is displayed during base driver installation even if user mode iWARP was not selected, as this option is applicable to Network Direct Kernel functionality as well.
4. If Windows Firewall is installed and active, select "Create an Intel® Ethernet iWARP Port Mapping Service rule in Windows Firewall" and the networks to which to apply the rule. If Windows Firewall is disabled or you are using a third party firewall, you will need to manually add this rule.
5. Continue with driver and software installation.

iWARP User Mode Installation on Microsoft Windows Server 2016 Nano Server

Follow the steps below to install the Intel® Ethernet User Mode iWARP Provider on Microsoft Windows Server 2016 Nano Server.

1. Create a directory from which to install the iWARP files. For example, C:\Nano\iwarp.
2. Copy the following files into your new directory:
    - \Disk\APPS\PROSETDX\Winx64\DRIVERS\i40wb.dll
    - \Disk\APPS\PROSETDX\Winx64\DRIVERS\i40wbmsg.dll
    - \Disk\APPS\PROSETDX\Winx64\DRIVERS\indv2.cat
    - \Disk\APPS\PROSETDX\Winx64\DRIVERS\indv2.inf
    - \Disk\APPS\PROSETDX\Winx64\DRIVERS\indv2.sys
3. Run the DISM command to inject the iWARP files into your Nano Server image, using the directory you created in step 1 for the AddDriver path parameter. For example, "DISM .../Add-Driver C:\Nano\iwarp"
4. Create an inbound firewall rule for UDP port 3935.
5. If desired, use the Windows PowerShell commands below to enable iWARP routing across IP Subnets.
    - Set-NetOffloadGlobalSetting -NetworkDirectAcrossIPSubnets Allow
    - Disable Adapter
    - Enable Adapter

Installing the Adapter

Select the Correct Slot

One open PCI-Express slot, x4, x8, or x16, depending on your adapter.

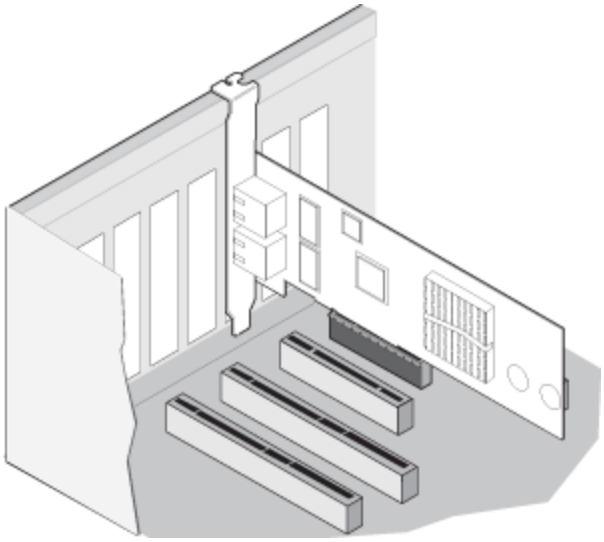
NOTE: Some systems have physical x8 PCI Express slots that actually only support lower speeds. Please check your system manual to identify the slot.

Insert the Adapter into the Computer

1. If your computer supports PCI Hot Plug, see your computer documentation for special installation instructions.
2. Turn off and unplug your computer. Then remove the cover.

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

3. Remove the cover bracket from an available slot.
4. Insert the adapter, pushing it 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.

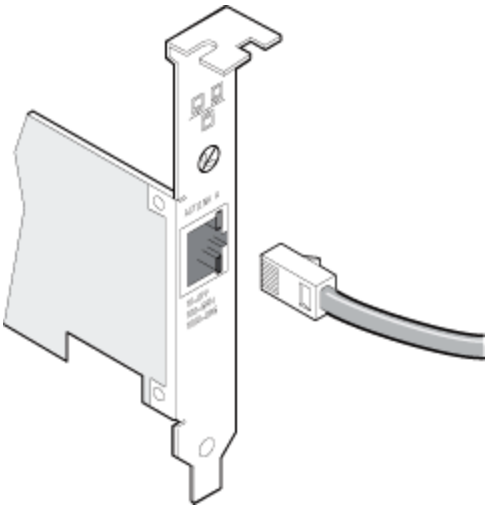
5. Secure the adapter bracket with a screw, if required.
6. Replace the computer cover and plug in the power cord.
7. Power on the computer.

Connecting Network Cables

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

Connect the RJ-45 Network Cable

Connect the RJ-45 network cable as shown:



The following table shows the maximum lengths for each cable type at a given transmission speed.

	Category 5	Category 6	Category 6a	Category 7
1 Gbps	100m	100m	100m	100m
10 Gbps	NA	55m	100m	100m
25 Gbps	NA	NA	NA	50m
40 Gbps	NA	NA	NA	50m

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.

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.

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.

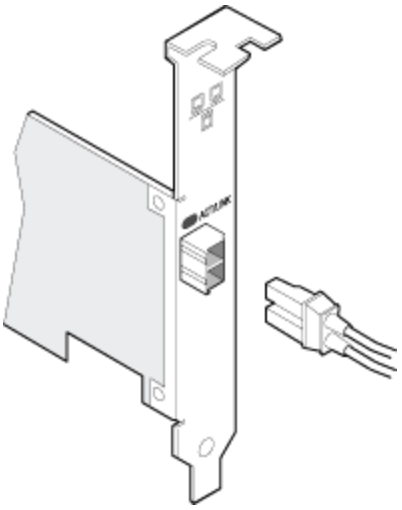
Remove and save the fiber optic connector cover. Insert a fiber optic cable into the ports on the network adapter bracket as shown below.

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 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

Intel Ethernet LR Server Adapters

Laser wavelength: 1310 nanometer (not visible)

Connector type: LC

Cable type: Single-mode fiber with 9.0µm core diameter

- Maximum cable length: 10 kilometers

Intel Ethernet 10 Gigabit SR Server Adapters

Laser wavelength: 850 nanometer (not visible)

Connector type: LC or SC

Cable type: Multi-mode fiber with 62.5µm core diameter

- Maximum cable length: 33 meters

Cable type: Multi-mode fiber with 50µm core diameter

- Maximum cable length: 300 meters

Intel Ethernet Gigabit SR Server Adapters

Laser wavelength: 850 nanometer (not visible)

Connector type: LC or SC

Cable type: Multi-mode fiber with 62.5µm core diameter

- Maximum cable length: 275 meters

Cable type: Multi-mode fiber with 50µm core diameter

- Maximum cable length: 550 meters

Supported SFP+ and QSFP+ Modules

Adapters Based on the 710 Series of Controllers

For information on supported media, see the following link:

http://www.intel.com/content/dam/www/public/us/en/documents/release-notes/xl710-ethernet-controller-feature-matrix.pdf

NOTES:

- Some Intel branded network adapters based on the X710/XL710 controller only support Intel branded modules. On these adapters, other modules are not supported and will not function.
- For connections based on the 710 series of controllers, support is dependent on your system board. Please see your vendor for details.
- In all cases Intel recommends using Intel optics; other modules may function but are not validated by Intel. Contact Intel for supported media types.
- In systems that do not have adequate airflow to cool the adapter and optical modules, you must use high temperature optical modules.
- For XXV710 based SFP+ adapters Intel recommends using Intel optics and cables. Other modules may function but are not validated by Intel. Contact Intel for supported media types.

82599-Based Adapters

NOTES:

- If your 82599-based Intel® Network Adapter came with Intel optics, or is an Intel® Ethernet Server Adapter X520-2, then it only supports Intel optics and/or the direct attach cables listed below.
- 82599-Based adapters support all passive and active limiting direct attach cables that comply with SFF-8431 v4.1 and SFF-8472 v10.4 specifications.

Supplier	Type	Part Numbers
SR Modules		

Intel	DUAL RATE 1G/10G SFP+ SR (bailed)	AFBR-703SDZ-IN2
Intel	DUAL RATE 1G/10G SFP+ SR (bailed)	FTLX8571D3BCV-IT
Intel	DUAL RATE 1G/10G SFP+ SR (bailed)	AFBR-703SDDZ-IN1
LR Modules		
Intel	DUAL RATE 1G/10G SFP+ LR (bailed)	FTLX1471D3BCV-IT
Intel	DUAL RATE 1G/10G SFP+ LR (bailed)	AFCT-701SDZ-IN2
Intel	DUAL RATE 1G/10G SFP+ LR (bailed)	AFCT-701SDDZ-IN1
QSFP Modules		
Intel	TRIPLE RATE 1G/10G/40G QSFP+ SR (bailed) (40G not supported on 82599)	E40GQSFPSR

The following is a list of 3rd party SFP+ modules that have received some testing. Not all modules are applicable to all devices.

Supplier	**Type**	**Part Numbers**
Finisar	SFP+ SR bailed, 10G single rate	FTLX8571D3BCL
Avago	SFP+ SR bailed, 10G single rate	AFBR-700SDZ
Finisar	SFP+ LR bailed, 10G single rate	FTLX1471D3BCL
Finisar	DUAL RATE 1G/10G SFP+ SR (No Bail)	FTLX8571D3QCV-IT
Avago	DUAL RATE 1G/10G SFP+ SR (No Bail)	AFBR-703SDZ-IN1
Finisar	DUAL RATE 1G/10G SFP+ LR (No Bail)	FTLX1471D3QCV-IT
Avago	DUAL RATE 1G/10G SFP+ LR (No Bail)	AFCT-701SDZ-IN1

Finisar	1000BASE-T SFP	FCLF8522P2BTL
Avago	1000BASE-T SFP	ABCU-5710RZ
HP	1000BASE-SX SFP	453153-001

82598-Based Adapters

NOTES:

- Intel® Network Adapters that support removable optical modules only support their original module type (i.e., the Intel® 10 Gigabit SR Dual Port Express Module only supports SR optical modules). If you plug in a different type of module, the driver will not load.
- 82598-Based adapters support all passive direct attach cables that comply with SFF-8431 v4.1 and SFF-8472 v10.4 specifications. Active direct attach cables are not supported.
- Hot Swapping/hot plugging optical modules is not supported.
- Only single speed, 10 Gigabit modules are supported.
- LAN on Motherboard (LOMs) may support DA, SR, or LR modules. Other module types are not supported. Please see your system documentation for details.

The following is a list of SFP+ modules and direct attach cables that have received some testing. Not all modules are applicable to all devices.

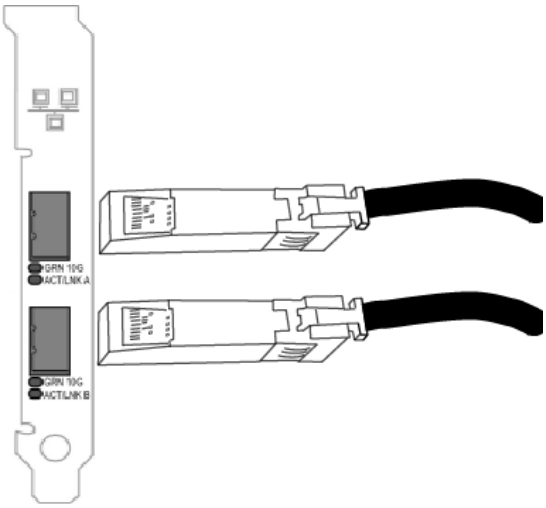
Supplier	Type	Part Numbers
Finisar	SFP+ SR bailed, 10G single rate	FTLX8571D3BCL
Avago	SFP+ SR bailed, 10G single rate	AFBR-700SDZ
Finisar	SFP+ LR bailed, 10G single rate	FTLX1471D3BCL
Molex	1m - Twin-ax cable	74752-1101
Molex	3m - Twin-ax cable	74752-2301
Molex	5m - Twin-ax cable	74752-3501
Molex	10m - Twin-ax cable	74752-9004
Tyco	1m - Twin-ax cable	2032237-2

Tyco	3m - Twin-ax cable	2032237-4
Tyco	5m - Twin-ax cable	2032237-6
Tyco	10m - Twin-ax cable	1-2032237-1

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 Attach Cable

Insert the Direct Attach network cable as shown below.



Type of cabling:

- 40 Gigabit Ethernet over SFP+ Direct Attached Cable (Twinaxial)
  - Length is 7 meters max.
- 25 Gigabit Ethernet over SFP28 Direct Attached Cable (Twinaxial)
  - Length is 5 meters max.
  - For optimal performance, must use CA-25G-L with RS-FEC and 25GBASE-CR
- 10 Gigabit Ethernet over SFP+ Direct Attached Cable (Twinaxial)
  - Length is 10 meters max.

PCI Hot Plug Support

Most Intel® Ethernet Server Adapters are enabled for use in selected servers equipped with Hot Plug support. Exceptions: Intel Gigabit Quad Port Server adapters do not support Hot Plug operations.

If you replace an adapter in a Hot Plug slot, do not place the removed adapter back into the same network until the server has rebooted (unless you return it to the same slot and same team as before). This prevents a conflict in having two of the same Ethernet addresses on the same network.

The system will require a reboot if you

- Change the primary adapter designator.
- Add a new adapter to an existing team and make the new adapter the primary adapter.
- Remove the primary adapter from the system and replace it with a different type of adapter.

NOTE: To replace an existing SLA-teamed adapter in a Hot Plug slot, first unplug the adapter cable. When the adapter is replaced, reconnect the cable.

PCI Hot Plug Support for Microsoft* Windows* Operating Systems

Intel® network adapters are enabled for use in selected servers equipped with PCI Hot Plug support and running Microsoft* Windows* operating systems. For more information on setting up and using PCI Hot Plug support in your server, see your hardware and/or Hot Plug support documentation for details. PCI Hot Plug only works when you hot plug an identical Intel network adapter.

NOTES:

- The MAC address and driver from the removed adapter will be used by the replacement adapter unless you remove the adapter from the team and add it back in. If you do not remove and restore the replacement adapter from the team, and the original adapter is used elsewhere on your network, a MAC address conflict will occur.
- For SLA teams, ensure that the replacement NIC is a member of the team before connecting it to the switch.

Microsoft* Windows* Driver and Software Installation and Configuration

Installing Windows Drivers and Software

NOTE: To successfully install or uninstall the drivers or software, you must have administrative privileges on the computer completing installation.

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 Device Manager.

NOTES:

- You must have administrator rights to install or use Intel PROSet for Windows Device Manager.
- Intel PROSet for Windows Device Manager and the IntelNetCmdlets module for Windows PowerShell* require the latest driver and software package for your Intel Ethernet devices. Please download the most recent driver and software package for your operating system from www.intel.com.
- On recent operating systems, older hardware may not support Intel PROSet for Windows Device Manager and the IntelNetCmdlets module for Windows PowerShell. In this case, the Intel PROSet tabs may not be displayed in the Windows Device Manager user interface, and the IntelNetCmdlets may display an error message stating that the device does not have an Intel driver installed.

Install the Drivers

NOTES:

- This will update the drivers for all supported Intel® network adapters in your system.
- 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 the support website.

Before installing or updating the drivers, insert your adapter(s) in the computer and plug in the network cable. When Windows discovers the new adapter, it attempts to find an acceptable Windows driver already installed with the operating system.

If found, the driver is installed without any user intervention. If Windows cannot find the driver, the Found New Hardware Wizard window is displayed.

Regardless of whether Windows finds the driver, it is recommended that you follow the procedure below to install the driver. Drivers for all Intel adapters supported by this software release are installed.

1. Download the latest drivers from the support website and transfer them to the system.
2. If the Found New Hardware Wizard screen is displayed, click **Cancel**.

3. Start the autorun located in the downloaded the software package. The autorun may automatically start after you have extracted the files.
4. Click **Install Drivers and Software**.
5. Follow the instructions in the install wizard.

NOTE: Intel® PROSet is installed by default when you install the device drivers.

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.

NOTES:

- You must install Intel® PROSet for Windows Device Manager if you want to use Intel® ANS teams or VLANs.
- Intel PROSet for Windows Device Manager is installed by default when you install the device drivers. For information on usage, see Using 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.

Installing the Base Driver and Intel® PROSet on Nano Server

Driver Installation

NOTE: Installing drivers requires administrator rights to the operating system.

To install drivers on Microsoft* Windows Server* Nano Server:

1. Identify which drivers to inject into the operating system.
2. Create a directory from which to install the drivers. For example, C:\Nano\Drivers.
3. Copy the appropriate drivers for the operating system and hardware. For example, "copy D:\PROXGB\Winx64\NDIS65*.* c:\Nano\Drivers /y"
4. If you are using the New-NanoServerImage module, use the above path for the -DriversPath parameter. For example, "New-NanoServerImage ...-DriversPath C:\Nano\Drivers"
5. If you are using DISM.exe as well, use the above path for the /AddDriver parameter. For example, "DISM .../Add-Driver C:\Nano\Drivers"

Intel PROSet Installation

To install Intel PROSet on Microsoft* Windows Server* Nano Server:

1. Use the New-NanoServerImage cmdlet to add the PROSetNS.zip file from the .\Disk\APPS\PROSETDX\NanoServer directory to your -CopyPath parameter.
2. Append the NanoSetup.ps1 file (located in the same directory) to your -SetupCompleteCommands parameter.

For example:

```
New-NanoServerImage ...
-CopyPath "<PATH>\PROSetNS.zip", "<PATH>\NanoSetup.ps1" `
-SetupCompleteCommands "PowerShell ""C:\NanoSetup.ps1"""
```

See the link below for more information on deploying a Nano Server image and using the cmdlet:
https://msdn.microsoft.com/en-us/library/mt126167.aspx

Command Line Installation for Base Drivers and Intel® PROSet

Driver Installation

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

NOTES:

- Intel® 10GbE Network Adapters do not support unattended driver installation.
- Intel PROSet cannot be installed with msiexec.exe. You must use DxSetup.exe.

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

DxSetup.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.

DxSetup.exe supports the following command line parameters:

Parameter	Definition
BD	Base Driver "0", do not install the base driver.

Parameter	Definition
	"1", install the base driver.
ANS	Advanced Network Services "0", do not install ANS. 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. 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.
FCOE	Fibre Channel over Ethernet "0", do not install FCoE. 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. If iSCSI is already installed, it will be uninstalled. "1", install FCoE. The iSCSI property requires DMIX=1.
IWARP_ ROUTING	iWARP routing "0", do not install iWARP routing. "1", install iWARP routing.
IWARP_ FIREWALL	Installs the iWARP firewall rule. For more information see iWARP (Internet Wide Area RDMA Protocol) section. "0", do not install iWARP firewall rule. "1", install iWARP firewall rule. If "1" is selected, the following parameters are allowed in addition to IWARP_FIREWALL. • IWARP_FIREWALL_DOMAIN [0|1] - Applies firewall rule to corporate domains. • IWARP_FIREWALL_PUBLIC [0|1] - Applies firewall rule to public networks

<table>
<tr><th>Parameter</th><th>Definition</th></tr>
<tr><td></td><td>• IWARP_FIREWALL_PRIVATE [0|1] - Applies firewall rule to private networks</td></tr>
<tr><td>FORCE</td><td>"0", check that the installed device supports a feature (FCOE, iSCSI) and only install the feature if such a device is found.
"1", install the specified features regardless of the presence of supporting devices.</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 PROSet 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>/uninstall
/x</td><td>Uninstall Intel PROSet and drivers.</td></tr>
</table>

NOTES:

- You must include a space between parameters.
- If you specify a path for the log file, the path must exist. If you do not specify a complete path, the install log will be created in the current directory.
- You do not need to specify default values. To install the base drivers, Intel PROSet, and ANS, the following examples are equivalent:

  ```
  DxSetup.exe
  ```

  ```
  DxSetup.exe BD=1 DMIX=1 ANS=1
  ```
- The ANS property should only be set to ANS=1 if DMIX=1 is set. If DMIX=0 and ANS=1, the ANS=1 is ignored and only the base driver will be installed.
- Even if FCOE=1 is passed, FCoE using DCB will not be installed if the operating system and installed adapters do not support it. If FORCE=1 is also passed, FCoE will be installed if the operating system supports it.
- Even if ISCSI=1 is passed, iSCSI using DCB will not be installed if the operating system and installed adapters do not support it. If FORCE=1 is also passed, iSCSI will be installed if the operating system supports it.
- Public properties are not case sensitive. No white space is allowed between characters. For example:

  ```
  DxSetup.exe /qn DMIX=1
  ```

  Any white space in "DMIX=1" makes the setting invalid.

Command line install examples

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

1. How to install the base driver:

   ```
   D:\DxSetup.exe DMIX=0 ANS=0
   ```
2. How to install the base driver using the logging option:

   ```
   D:\DxSetup.exe /l C:\installBD.log DMIX=0 ANS=0
   ```
3. How to install Intel PROSet and ANS silently:

   ```
   D:\DxSetup.exe DMIX=1 ANS=1 /qn
   ```
4. How to install Intel PROSet without ANS silently:

   ```
   D:\DxSetup.exe DMIX=1 ANS=0 /qn
   ```
5. How to install components but deselect ANS:

   ```
   D:\DxSetup.exe DMIX=1 ANS=0 /qn /liew C:\install.log
   ```
   The /liew log option provides a log file for the Intel PROSet 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:\DxSetup.exe ANS=1`.

Modify and Upgrade

You can use DxSetup.exeto modify or upgrade your drivers and software. If a feature is already installed, the public property for that feature will default to 1 and if a feature is not installed, the public property for that feature will default to 0. Running DxSetup.exewithout specifying properties will upgrade all installed software. You can remove installed software (except for base drivers) by setting the property to 0. If you uninstall PROSet (DMIX=0), all features that rely on PROSet will also be removed.

Windows Server Core

Command Line Options

SetupBD.exe supports the following command line switches.

NOTE: You must include a space between switches.

Switch	Description
/s	silent install
/r	force reboot (must be used with the /s switch)
/nr	no reboot (must be used with the /s switch. This switch is ignored if it is included with the /r switch)

Examples:

Option	Description
SetupBD	Installs and/or updates the driver(s) and displays the GUI.
SetupBD /s	Installs and/or updates the driver(s) silently.
SetupBD /s /r	Installs and/or updates the driver(s) silently and forces a reboot.
SetupBD /s /r /nr	Installs and/or updates the driver(s) silently and forces a reboot (/nr is ignored).

Other information

You can use the **/r** and **/nr** switches only with a silent install (i.e. with the "/s" option).

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:

- You must have Administrator privileges to run scripts. If you do not have Administrator privileges, you will not receive an error, the script just will not run.
- 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.

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\Wired Networking\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\Wired Networking\DMIX).
3. Type the following:
   ```
   SaveRestore.ps1 -Action Restore -ConfigPath e:\settings.txt
   ```

Configuring Device Features

Configuring with Intel® PROSet for Windows* Device Manager

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 Device Manager.

NOTES:

- You must have administrator rights to install or use Intel PROSet for Windows Device Manager.
- Intel PROSet for Windows Device Manager and the IntelNetCmdlets module for Windows PowerShell* require the latest driver and software package for your Intel Ethernet devices. Please download the most recent driver and software package for your operating system from www.intel.com.
- On recent operating systems, older hardware may not support Intel PROSet for Windows Device Manager and the IntelNetCmdlets module for Windows PowerShell. In this case, the Intel PROSet tabs may not be displayed in the Windows Device Manager user interface, and the IntelNetCmdlets may display an error message stating that the device does not have an Intel driver installed.

Changing Intel PROSet Settings Under Windows Server Core

You can use the Intel NetCmdlets for Microsoft* Windows PowerShell* to change most Intel PROSet settings under Windows Server Core. Please refer to the aboutIntelNetCmdlets.hlp.txt help file.

For iSCSI Crash Dump configuration, use the Intel NetCmdlets for Microsoft* Windows PowerShell* and refer to the aboutIntelNetCmdlets.help.txt help file.

NOTE: Support for the Intel PROSet command line utilities (prosetcl.exe and crashdmp.exe) has been removed, and is no longer installed. This functionality has been replaced by the Intel Netcmdlets for Microsoft* Windows PowerShell*. Please transition all of your scripts and processes to use the Intel Netcmdlets for Microsoft Windows PowerShell.

Compatibility Notes

The following devices do not support Intel PROSet for Windows Device Manager.

- Intel® 82552 10/100 Network Connection
- Intel® 82567V-3 Gigabit Network Connection
- Intel® X552 10G Ethernet devices
- Intel® X553 10G Ethernet devices
- Any platform with a System on a Chip (SoC) processor that includes either a server controller (designated by an initial X, such as X552) or both a server and client controller (designated by an initial I, such as I218)
- Devices based on the Intel® Ethernet Controller X722

Link Speed tab

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

Advanced Tab

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.

VLANs Tab

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.

Teaming Tab

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.

Power Management Tab

The Intel® PROSet **Power Management** tab replaces the standard Microsoft* Windows* Power Management tab in Device Manager. The standard Windows power management functionality is included on the Intel PROSet tab.

NOTES:

- The options available on the Power Management tab are adapter and system dependent. Not all adapters will display all options. 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).
- 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.

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.

NOTE To use the Minihelp property for any cmdlet in the module, append "| Select Minihelp" to the full cmdlet syntax. For example:

```
Get-IntelNetAdapterSetting - Name "Intel(R) Ethernet Server Adapter
x520-2" -RegistryKeyword *RSS | Select Minihelp
```

The IntelNetCmdlets module for Windows PowerShell is installed by default when you install Intel PROSet for Windows Device Manager. After installation, 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. For example:

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

NOTE: If you include a trailing backslash ("\") at the end of the Import-Module command, the import operation will fail. In Microsoft Windows* 10 and Windows Server* 2016, the auto-complete function appends a trailing backslash. If you use auto-complete when entering the Import-Module command, delete the trailing backslash from the path before pressing Return to execute the command.

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.

NOTES:

- 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.
- The Get-IntelNetAdapterStatus -Status General cmdlet may report the status "Link Up - This device is not linked at its maximum capable speed". In that case, if your device is set to auto-negotiate, you can adjust the speed of the device's link partner to the device's maximum speed. If the device is not set to auto-negotiate, you can adjust the device's speed manually, but you must ensure the link partner is set at the same speed.

Configuring with Intel® PROSet Adapter Configuration Utility

The Intel® PROSet Adapter Configuration Utility, or Intel® PACU, is a graphical user interface which allows you to configure and manage supported Intel® Ethernet Adapters (see Compatibility Notes below).

NOTE: The Intel® PROSet Adapter Configuration Utility is available only on Microsoft* Windows Server* 2019 and later.

The general configuration steps are as follows:

1. Select an adapter in the Adapter Selection panel.
2. Select a setting to configure from the Adapter Settings panel.
3. Select or enter the desired value(s) for the selected setting.
4. Click the "Apply Changes" button.

Compatibility Notes

The following devices do not support Intel PROSet Adapter Configuration Utility.

- Intel® 82552 10/100 Network Connection
- Intel® 82567V-3 Gigabit Network Connection
- Intel® X552 10G Ethernet devices
- Intel® X553 10G Ethernet devices
- Any platform with a System on a Chip (SoC) processor that includes either a server controller (designated by an initial X, such as X552) or both a server and client controller (designated by an initial I, such as I218)
- Devices based on the Intel® Ethernet Controller X722

Device Features

Adapter Teaming

Intel® Advanced Network Services (Intel® ANS) Teaming 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 creating a team or adding team members, make sure each team member has been configured similarly. Settings to check include VLANs and QoS Packet Tagging, Jumbo Packets, and the various offloads. Pay particular attention when using different adapter models or adapter versions, as adapter capabilities vary.

Configuration Notes

- You must install the latest Microsoft* Windows* 10 updates before you can create Intel ANS Teams or VLANs on Windows 10 systems. Any Intel ANS Teams or VLANs created with a previous software/driver release on a Windows 10 system will be corrupted and cannot be upgraded. The installer will remove these existing teams and VLANs.
    - If you are running Windows 10 Anniversary edition (RS1) you will need to install Intel LAN software v22.1 or newer.
    - If you are running Windows 10 Creators Update (RS2) you will need to install Intel LAN software v22.3 or newer.

      NOTE: If you are running Windows 10 Creators Update (RS2) and are using Intel LAN software release v22.1 or v22.2, then ANS will not function until you update to the Intel LAN software v22.3 or newer.

- Microsoft* Windows Server* 2012 R2 is the last Windows Server operating system version that supports Intel Advanced Networking Services (Intel ANS). Intel ANS is not support on Microsoft Windows Server 2016 and later.
- 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.
- Not all team types are available on all operating systems.
- Be sure to use the latest available drivers on all adapters.
- Not all Intel devices support Intel ANS or Intel PROSet. Intel adapters that do not support Intel ANS or 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.
- You cannot create a team that includes both Intel X710/XL710-based devices and Intel® I350-based devices. These devices are incompatible together in a team and will be blocked during team setup. Previously created teams that include this combination of devices will be removed upon upgrading.
- NDIS 6.2 introduced new RSS data structures and interfaces. Because of 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 a team is bound to a Hyper-V virtual NIC, you cannot change the Primary or Secondary adapter.
- To assure a common feature set, some advanced features, including hardware offloading, are automatically disabled when an adapter that does not support Intel PROSet is added to a team.
- Hot Plug operations in a Multi-Vendor Team may cause system instability. We recommended that you restart the system or reload the team after performing Hot Plug operations with a Multi-Vendor Team.
- Spanning tree protocol (STP) should be disabled on switch ports connected to teamed adapters in order to prevent data loss when the primary adapter is returned to service (failback). 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/Data Center Bridging will be automatically disabled when an adapter is added to a team with non-FCoE/DCB capable adapters.

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.

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**.
10. Click **Finish**.

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

To change this setting in Windows PowerShell, use the New-IntelNetTeam cmdlet. For example, New-IntelNetTeam -TeamMemberNames "<adapter1_name>", "<adapter2_name>" -TeamMode AdapterFaultTolerance -TeamName "<team_name>".

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.

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**.

To change this setting in Windows PowerShell, use the Add-IntelNetTeamMember cmdlet. For example, Add-IntelNetTeamMember -TeamName "<team_name>" -Name "<adapter_name>".

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**.

To change this setting in Windows PowerShell, use the Set-IntelNetTeam cmdlet. For example, Set-IntelNetTeam -TeamName "<team_name>" -NewTeamName "<new_team_name>".

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.

To change this setting in Windows PowerShell, use the Remove-IntelNetTeamMember cmdlet. For example, Remove-IntelNetTeamMember -Name "<adapter_name>".

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.

Teaming and VLAN Considerations When Replacing Adapters

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.

To do this in Windows PowerShell, use the cmdlets above to remove the team member or team, then add the replaced adapter to the team (or create a new team).

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

NOTES:

- 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.
- Data Center Bridging (DCB) is incompatible with Microsoft Server LBFO Teams. Do not create an LBFO team using Intel 10G ports when DCB is installed. Do not install DCB if Intel 10G ports are part of an LBFO team. Install failures and persistent link loss may occur if DCB and LBFO are used on the same port.

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 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 device does not support the desired team type or does not support teaming at all.
- The operating system does not support the desired team type.
- 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.
- TOE (TCP Offload Engine) enabled devices cannot be added to an ANS team and will not appear in the list of available adapters.
- 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” or "Enabled" 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.

Teaming Modes

Adapter Fault Tolerance (AFT) - provides automatic redundancy for a server's network connection. If the primary adapter fails, the secondary adapter takes over. Adapter Fault Tolerance supports two to eight adapters per team. This teaming type works with any hub or switch. All team members must be connected to the same subnet.

- Switch Fault Tolerance (SFT) - provides failover between two adapters connected to separate switches. Switch Fault Tolerance supports two adapters per team. Spanning Tree Protocol (STP) must be enabled on the switch when you create an SFT team. When SFT teams are created, the Activation Delay is automatically set to 60 seconds. This teaming type works with any switch or hub. All team members must be connected to the same subnet.
- Adaptive Load Balancing (ALB) - provides load balancing of transmit traffic and adapter fault tolerance. In Microsoft* Windows* operating systems, you can also enable or disable receive load balancing (RLB) in ALB teams (by default, RLB is enabled).
- 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. This teaming type works with any switch.
- Static Link Aggregation (SLA) - provides increased transmission and reception throughput in a team of two to eight adapters. This team type replaces the following team types from prior software releases: Fast EtherChannel*/Link Aggregation (FEC) and Gigabit EtherChannel*/Link Aggregation (GEC). This type also includes adapter fault tolerance and load balancing (only routed protocols). This teaming type requires a switch with Intel Link Aggregation, Cisco* FEC or GEC, or IEEE 802.3ad Static Link Aggregation capability.

  All adapters in a Link Aggregation team running in static mode must run at the same speed and must be connected to a Static Link Aggregation capable switch. If the speed capability of adapters in a Static Link Aggregation team are different, the speed of the team is dependent on the lowest common denominator.

- IEEE 802.3ad Dynamic Link Aggregation - creates one or more teams using Dynamic Link Aggregation with mixed-speed adapters. Like the Static Link Aggregation teams, Dynamic 802.3ad teams increase transmission and reception throughput and provide fault tolerance. This teaming type requires a switch that fully supports the IEEE 802.3ad standard.
- Multi-Vendor Teaming (MVT) - adds the capability to include adapters from selected other vendors in a team. If you are using a Windows-based computer, you can team adapters that appear in the Intel PROSet teaming wizard.

IMPORTANT:

- Be sure to use the latest available drivers on all adapters.
- 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 Frames, and the various off-loads. 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:
  - 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.

NOTES:

- Hot Plug operations for an adapter that is part of a team are only available in Windows Server.
- For SLA teams, all team members must be connected to the same switch. For AFT, ALB, and RLB teams, all team members must belong to the same subnet. The members of an SFT team must be connected to a different switch.
- Teaming only one adapter port is possible, but provides no benefit.

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 allback 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**.
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.

To do this in Windows PowerShell, use the Set-IntelNetTeam cmdlet. For example, Set-IntelNetTeam -TeamName "Team 1" -PrimaryAdapterName "<adapter1_name>" -SecondaryAdapterName "<adapter2_name>".

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 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.

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.
- You cannot use and Intel AMT enabled adapter 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

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).
- 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.

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.

To remove a phantom team in Windows PowerShell, use the Remove-IntelNetTeamMember cmdlet. To remove a phantom VLAN in Windows PowerShell, use the Remove-IntelNetVLAN cmdlet. For example, Remove-IntelNetVLAN -ParentName "<adapter_name>" -VLANID "1" .

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.

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.

Range	• Enabled • Disabled

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Adaptive Inter-Frame Spacing" –DisplayValue "Enabled"

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.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "DMA Coalescing" –DisplayValue "Enabled"

Forward Error Correction (FEC) Mode

Allows you to set the Forward Error Correction (FEC) mode. FEC improves link stability, but increases latency. Many high quality optics, direct attach cables, and backplane channels provide a stable link without FEC.

The driver allows you to set the following FEC Modes:

- Auto FEC - Sets the FEC Mode based on the capabilities of the attached cable.
- CL108 RS-FEC - Selects only RS-FEC ability and request capabilities.
- CL74 FC-FEC/BASE-R - Selects only BASE-R ability and request capabilities.
- No FEC - Disables FEC.

NOTE: For devices to benefit from this feature, link partners must have FEC enabled.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "FEC Mode" –DisplayValue "Auto FEC"

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.

NOTE:

- For adapters to benefit from this feature, link partners must support flow control frames.

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

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Flow Control" –DisplayValue "Auto Negotiation"

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.

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.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Gigabit Master Slave Mode" –DisplayValue "Auto Detect"

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.

Range	• Adaptive

	• Extreme • High • Medium • Low • Minimal • Off

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Interrupt Moderation Rate" –DisplayValue "Adaptive"

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.

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

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "IPv4 Checksum Offload" –DisplayValue "Tx Enabled"

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.

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.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Jumbo Packet" –DisplayValue "4088 Bytes"

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.

Range	• Enabled • Disabled

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Large Send Offload V2 (IPv4)" –DisplayValue "Enabled"

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.

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.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Locally Administered Address" –DisplayValue "<desired address>"

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.

Range	Enabled, Disabled

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Log Link State Event" –DisplayValue "Enabled"

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.

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

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Low Latency Interrupts" –DisplayValue "Port-Based"

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.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "NVGRE Encapsulated Task Offload" –DisplayValue "Enabled"

Performance Options

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.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Profile" –DisplayValue "Standard Server"

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.

Power 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.

On the device's Device Manager property sheet, this setting is found on the Power Management tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Wake on Link Settings" –DisplayValue "Enabled"

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.

Range	The range varies with the operating system and adapter.

Ultra Low Power Mode When Cable is Disconnected

Enabling Ultra Low Power (ULP) mode significantly reduces power consumption when the network cable is disconnected from the device.

NOTE: If you experience link issues when two ULP-capable devices are connected back to back, disable ULP mode on one of the devices.

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 on any 10 GbE, 25 GbE, or 40 GbE adapter.

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.

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.

NOTES:

- To use the Wake on Directed Packet feature, WoL must first be enabled in the EEPROM using BootUtil.
- If **Reduce speed during standby** is enabled, then **Wake on Magic Packet** and/or **Wake on directed packet** must be enabled. If both of these options are disabled, power is removed from the adapter during standby.
- **Wake on Magic Packet from power off state** has no effect on this option.

WoL Supported Devices

The following adapters support WoL only on Port A:

- Intel® Ethernet Server Adapter I350-T2
- Intel® Ethernet Server Adapter I350-T4
- Intel® Ethernet Server Adapter I340-T2
- Intel® Ethernet Server Adapter I340-T4
- Intel® Ethernet Server Adapter I340-F4
- Intel® Gigabit ET2 Quad Port Server Adapter
- Intel® PRO/1000 PF Quad Port Server Adapter
- Intel® PRO/1000 PT Quad Port LP Server Adapter
- Intel® PRO/1000 PT Quad Port Server Adapter
- Intel® PRO/1000 PT Dual Port Network Connection
- Intel® PRO/1000 PT Dual Port Server Connection
- Intel® PRO/1000 PT Dual Port Server Adapter
- Intel® PRO/1000 PF Dual Port Server Adapter
- Intel® Gigabit PT Quad Port Server ExpressModule

The following adapters do not support WoL:

- Intel® PRO/1000 MT Quad Port Server adapter
- Intel® Gigabit VT Quad Port Server Adapter
- Intel® Ethernet Server Adapter X520-2
- Intel® Ethernet Server Adapter X520-1
- Intel® Ethernet Server Adapter X540-T1
- Intel® Ethernet Converged Network Adapter X540-T2
- Intel® Ethernet Converged Network Adapter X540-T1
- Intel® Ethernet Converged Network Adapter X710-2
- Intel® Ethernet Converged Network Adapter X710-4
- Intel® Ethernet Converged Network Adapter X710-T4
- Intel® Ethernet Converged Network Adapter X710
- Intel® Ethernet Converged Network Adapter XL710-Q1
- Intel® Ethernet Converged Network Adapter XL710-Q2

Most Intel 10GbE Network Adapters do not support Wake on LAN on any port.

- The Intel® Ethernet Converged Network Adapter X550-T1 and Intel® Ethernet Converged Network Adapter X550-T2 have a manageability/AUX power connector. These devices only support WoL if AUX power is supplied via this connector. Note that this is system and adapter specific. Some devices with this connector do not support WoL. Some systems do not provide the correct power connection. See your system documentation for details.

The following 10GbE Network Adapters support Wake on LAN on all ports

- Intel® Ethernet Server Adapter X550-T2 for OCP
- Intel® Ethernet Server Adapter X550-T1 for OCP

Priority & VLAN Tagging

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

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

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To set this in Windows Powershell, first disable DCB, then set priority and VLAN tagging. For example:

Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "DCB" -DisplayValue "Disabled"

Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Packet Priority & VLAN" -DisplayValue "VLAN Enabled"

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.

Tagging is enabled and disabled in Microsoft* Windows* Server* using the "QoS Packet Tagging" field in the **Advanced** tab in Intel® PROSet. For other versions of the Windows operating system, tagging is enabled using the "Priority/VLAN Tagging" setting on the **Advanced** tab in Intel® PROSet.

To set this in Windows Powershell, first disable DCB, then set QoS using the Priority and VLAN tagging DisplayName in the cmdlet. For example:

Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "DCB" -DisplayValue "Disabled"

Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Packet Priority & VLAN" -DisplayValue "VLAN Enabled"

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

The supported operating systems, including Microsoft* Windows Server*, have a utility for 802.1p packet prioritization. For more information, see the Windows system help and Microsoft's knowledge base.

NOTE: The first generation Intel® PRO/1000 Gigabit Server Adapter (PWLA 8490) does not support QoS frame tagging.

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.

Range	128-4096, in intervals of 64, for all adapters.
Recommended Value	Teamed adapter: 256 Using IPSec and/or multiple features: 352

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Receive Buffers" –DisplayValue "256"

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.

  NOTES:

  - This setting only affects NUMA systems. It will have no effect on non-NUMA systems.
  - Choosing a value greater than the number of NUMA nodes present in the system selects the NUMA node closest to the device.

- **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).

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.

  On the device's Device Manager property sheet, this setting is found on the Advanced tab.

  To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Receive Side Scaling" –DisplayValue "Enabled"

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 is 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:

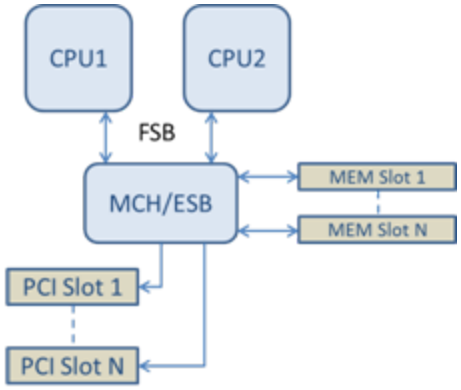
- FCoE NUMA Node Count = 2
- FCoE Starting NUMA Node = 0
- FCoE Starting Core Offset = 0

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:

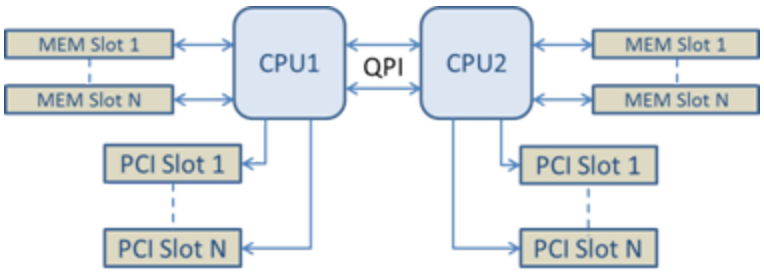
- FCoE NUMA Node Count = 1
- FCoE Starting NUMA Node = 2
- FCoE Starting Core Offset = 0

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.

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.

NOTE:

- Fiber-based adapters operate only in full duplex at their native speed.

The settings available when auto-negotiation is disabled are dependent on your device. Not all speeds are available on all devices. 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.
- The Link Speed tab may display a blue informational icon with a mouse-over message "This device is not linked at its maximum capable speed". In that case, if your device is set to auto-negotiate, you can adjust the speed of the device's link partner to the device's maximum speed. If the device is not set to auto-negotiate, you can adjust the device's speed manually, but you must ensure the link partner is set at the same speed.

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.

If the adapter cannot establish link with the gigabit link partner using auto-negotiation, set the adapter to **1 Gbps Full duplex**.

Intel 10 gigabit fiber-based adapters and SFP direct-attach devices operate only in full duplex, and only at their native speed. Multi-speed 10 gigabit SFP+ fiber modules support full duplex at 10 Gbps and 1 Gbps.

Auto-negotiation and Auto-Try are not supported on devices based on the Intel® Ethernet Connection X552 controller and Intel® Ethernet Connection X553 controller.

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.**

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.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Speed & Duplex" –DisplayValue "Auto Negotiation"

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.
- For best performance, on the host use 'Set-VMNetworkAdapter -IovQueuePairsRequested 4' on the VF to allow the virtual network to use 4 queues (maximum supported value) and assign 4 or more virtual CPUs to the connected VM. In the VM, set 'Maximum number of Receive Queues" in the VF's adapter properties to 4.
- SR-IOV is not supported with ANS teams.
- VMWare ESXi does not support SR-IOV on 1GbE ports.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "SR-IOV" –DisplayValue "Enabled"

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.

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

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "TCP Checksum Offload (IPv4)" –DisplayValue "Tx Enabled"

TCP/IP Offloading Options

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 Manager. 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.

To change this setting in Windows PowerShell use the Get-IntelNetAdapter cmdlet. For example: Get-IntelNetAdapter –Name "<adapter_name" | fl

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.

Range	128-16384, in intervals of 64, for 10 Gigabit Server Adapters. 128-4096, in intervals of 64, for all other adapters.

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Transmit Buffers" –DisplayValue "128"

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.

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

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "UDP Checksum Offload (IPv4)" –DisplayValue "Tx Enabled"

Virtual LANs

Overview

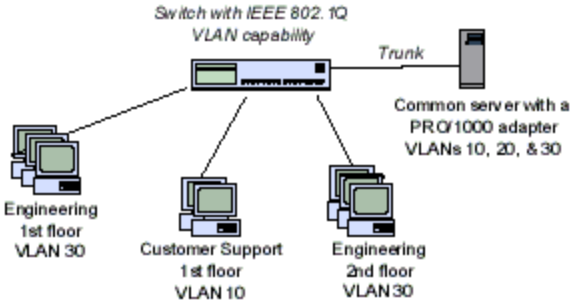
NOTES:

- You must install the latest Microsoft* Windows* 10 updates before you can create Intel ANS Teams or VLANs on Windows 10 systems. Any Intel ANS Teams or VLANs created with a previous software/driver release on a Windows 10 system will be corrupted and cannot be upgraded. The installer will remove these existing teams and VLANs.
  - If you are running Windows 10 Anniversary edition (RS1) you will need to install Intel LAN software v22.1 or newer.
  - If you are running Windows 10 Creators Update (RS2) you will need to install Intel LAN software v22.3 or newer.

    NOTE: If you are running Windows 10 Creators Update (RS2) and are using Intel LAN software release v22.1 or v22.2, then ANS will not function until you update to the Intel LAN software v22.3 or newer.

- Microsoft Windows Server 2012 R2 is the last Windows Server operating system version that supports Intel Advanced Networking Services (Intel ANS). Intel ANS is not supported on Microsoft Windows Server 2016 and later.
- 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.

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.
- Intel ANS VLANs are not compatible with the Microsoft Hyper-V virtual machine switch. If you want to bind the virtual machine switch to a VLAN, you must create the VLAN from within the Virtual Switch Manager.
- To set up IEEE VLAN membership (multiple VLANs), the adapter must be attached to a switch with IEEE 802.1Q VLAN capability.
- A maximum of 64 VLANs per network port or team are supported by Intel software.
- 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.
- The Intel PRO/100 VE and VM Desktop Adapters and Network Connections can be used in a switch based VLAN but do not support IEEE Tagging.
- 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.

CAUTION: When using IEEE 802 VLANs, settings must match between the switch and those adapters using the VLANs.

Configuring VLANs in Microsoft* Windows*

On the device's Device Manager property sheet, this setting is found on the VLANs tab.

To add a VLAN using Windows PowerShell, use the Add-IntelNetVLAN cmdlet. For example, Add-IntelNetVLAN -ParentName "Name" -VLANID "1".

To remove a VLAN using Windows PowerShell, use the Remove-IntelNetVLAN cmdlet. For example, Remove-IntelNetVLAN -ParentName "Name" -VLANID "1".

CAUTION:

- **VLANs cannot be used on teams that contain non-Intel network adapters**
- **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:

- The VLAN ID keyword is supported. The VLAN ID must match the VLAN ID configured on the switch. Adapters with VLANs must be connected to network devices that support IEEE 802.1Q.
- 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 Using Intel® Network Adapters in a Microsoft* Hyper-V* Environment.
- 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 Using Intel® Network Adapters in a Microsoft* Hyper-V* Environment.

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 be established 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".

Range	• On • Off • Auto Detect

On the device's Device Manager property sheet, this setting is found on the Advanced tab.

To change this setting in Windows PowerShell use the Set-IntelNetAdapterSetting cmdlet. For example: Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Wait for Link" –DisplayValue "Off"

Intel Network Drivers for DOS

The NDIS2 (DOS) driver is provided solely for the purpose of loading other operating systems -- for example, during RIS or unattended installations. It is not intended as a high-performance driver.

You can find adapter drivers, PROTOCOL.INI files, and NET.CFG files in the PRO100\DOS or PRO1000\DOS directory in the download folder. For additional unattended install information, see the text files in the operating system subdirectories under the APPS\SETUP\PUSH directory.

Automatic or Explicit Configuration of a Single NIC or Multiple NICs

When the driver finds that only one adapter is installed in the system, it will use that adapter regardless of whether or not parameters in PROTOCOL.INI are present or correct. If the parameters do not match the actual configuration, the driver will display warning messages indicating that the parameter was not used.

One instance of the driver must be loaded for each adapter that is activated. When multiple adapters are installed, the SLOT parameter becomes advisable but not required.

The determination as to which adapter each driver will control should be made by the user based on the protocol stack(s) bound to each driver, and based on the network that is connected to each adapter. The "BINDINGS" list in each protocol stack's PROTOCOL.INI section establishes the relationship between protocol stacks and drivers. The SLOT parameter in the driver's PROTOCOL.INI section establishes the relationship between drivers and adapters, and a value can be provided for each driver loaded. If a SLOT parameter is not specified, the first driver instance will load on the first NIC/Port found in the scanning list, the second driver instance will load on the second NIC/Port found in the scanning list, etc. When the driver detects multiple NICs/Ports it will report all of the possible slots. The only way for the driver to know which driver instance is being loaded is to use the DRIVERNAME parameter instance number. Therefore, it is essential that the DRIVERNAME parameter instance syntax defined below be used correctly.

The adapters are automatically configured by the PCI system BIOS when the system boots. The driver queries the PCI BIOS and obtains all of the adapter's configuration information. BIOS scanning using mechanisms 1 and 2, as defined in the PCI BIOS specification, are supported. The SLOT number is actually the encoded value of the PCI adapter's device location, which is defined as shown below. The SLOT value reported by the driver and entered by the user is the value of bits 0 through 15. In versions of the driver prior to 2.01, the SLOT value reported by the driver and entered by the user was shifted right by 3 bits (divided by 8) so that SLOT 0x0088 was actually entered into PROTOCOL.INI as 0x0011. This doesn't allow for multi-function devices to be specified with this SLOT parameter. So starting with v2.01, the driver does not shift the input parameter by 3 bits and SLOT 0x0088 would be entered as 0x0088. This also allows for specifying slot 0x0081 = Bus 0 Device 16 Function 1. If the driver finds that the entered SLOT number is not found in its slot list table, it may be because the SLOT uses the older convention (shifted right). The driver then tries to match this old style slot parameter to a slot in the slot list and loads on that slot if it finds a match. This is done for backward compatibility.

Configuring with the PROTOCOL.INI File

The configuration parameters listed below are supported through the PROTOCOL.INI file. When the machine has a single adapter, all the parameters (except DRIVERNAME) are optional; when the machine has multiple adapters, some of the parameters are required.

DRIVERNAME

This is the only parameter required for all configurations. This parameter is essentially an "instance ID". Each instance of the driver must create a unique instance name, both to satisfy DOS driver requirements, and to make it possible to find the parameters for the instance in the PROTOCOL.INI file.

When the driver initializes, it tries to find previously loaded instances of itself. If none are found, the driver calls itself "E1000$", and looks for that name in the PROTOCOL.INI file to find its parameters. If one or more instances are found, the driver calls itself "E100x$", where 'x' is one more than the value used by the most recently loaded instance. So, in this scenario, the second driver calls itself "E1002$", the third calls itself "E1003$", and so on; there is no driver called "E1001$". Up to 10 drivers can be loaded in a single system in this way.

Syntax:	DRIVERNAME = [E1000$ | E1002$ | etc.]
Example:	DRIVERNAME = E1000$
Default:	None, this is a required parameter.
Normal Behavior:	The driver finds its section in PROTOCOL.INI by matching its instance ID to the value for this parameter.
Possible Errors:	The device driver uses a DOS function to display the name of the driver it is expecting. This function cannot display a '$' character. For this reason, the user may see a message referring to this value without the '$'; the user must remember to enter the '$' character as part of the parameter's value.

SPEEDDUPLEX

The parameter disables Auto-Speed-Detect and causes the adapter to function at the speed indicated. Do not include this parameter if you want your Gigabit adapter to connect at 1000 Mbps.

Syntax:	SPEEDDUPLEX = [0 | 1 | 2 | 3]
Example:	SPEEDDUPLEX = 2
Default:	Parameter not included in PROTOCOL.INI
Normal Behavior:	0 = 10 Mbps half duplex 1 = 10 Mbps full duplex 2 = 100Mbps half duplex 3 = 100Mbps full duplex
Possible Errors:	If the SPEEDDUPLEX parameter is set to an invalid value: • The parameter is ignored and the default (Auto-Speed-Detect) is used • A message indicates a "Parameter value out of range" error

SLOT

This parameter makes it possible for the driver to uniquely identify which of the adapters is to be controlled by the driver. The parameter can be entered in hexadecimal or decimal.

Syntax:	SLOT = [0x0..0x1FFF] SLOT = [0..8191]
Examples:	SLOT = 0x1C SLOT = 28
Default:	The driver will Auto-Configure if possible.
Normal Behavior:	The driver uses the value of the parameter to decide which adapter to control.
Possible Errors:	If only one adapter is installed, and the value does not correctly indicate the adapter slot: • A message indicates that the value does not match the actual configuration • The driver finds the adapter and uses it If more than one adapter is installed, and the value does not correctly indicate an adapter slot: • A message indicates possible slots to use • The driver loads on the next available slot

NODE

This parameter sets the Individual Address of the adapter, overriding the value read from the EEPROM.

Syntax:	NODE = "12 hexadecimal digits" • The value must be exactly 12 hexadecimal digits, enclosed in double quotes. • The value can not be all zeros. • The value can not have the Multicast bit set (LSB of 2nd digit = 1).
Example:	NODE = “00AA00123456”
Default:	Value from EEPROM installed on adapter
Normal Behavior:	The Current Station Address in the NDIS MAC Service-Specific Characteristics (MSSC) table is assigned the value of this parameter. The adapter hardware is programmed to receive frames with the destination address equal to the Current Station Address in the MSSC table. The Permanent Station Address in the MSSC table will be set to reflect the node address read from the adapter's EEPROM.
Possible Errors:	If any of the rules described above are violated, the driver treats this as a fatal error and an error message occurs, indicating the correct rules for forming a proper address.

ADVERTISE

This parameter can be used to restrict the speeds and duplexes advertised to a link partner during auto-negotiation. If AutoNeg = 1, this value is used to determine what speed and duplex combinations are advertised to the link partner. This field is treated as a bit mask.

Syntax:	ADVERTISE = [1 | 2 | 4 | 8 | 0x20 | 0x2F]: 0x01 = 10 Half, 0x02 = 10 Full, 0x04 = 100 Half, 0x08 = 100 Full, 0x20 = 1000 Full, 0x2F = all rates
Example:	ADVERTISE = 1
Default:	0x2F (all rates are supported)
Normal Behavior:	By default all speed/duplex combinations are advertised.
Possible Errors:	An error message is displayed if the value given is out of range.

FLOWCONTROL

This parameter, which refers to IEEE 802.3x flow control, helps prevent packets from being dropped and can improve overall network performance. Specifically, the parameter determines what flow control capabilities the adapter advertises to its link partner when auto negotiation occurs. This setting does NOT force flow control to be used. It only affects the advertised capabilities.

NOTES:

- Due to errata in the 82542 silicon, the chip is not able to receive PAUSE frames if the ReportTxEarly parameter is set to 1. Thus, if ReportTxEarly =1 and the driver is running on an adapter using this silicon (such as the PWLA8490), the driver will modify the FlowControl parameter to disable the ability to receive PAUSE frames.
- If half-duplex is forced or auto-negotiated, the driver will completely disable flow control.

Syntax:	FLOWCONTROL = [0 | 1 | 2 | 3 |0xFF]
Example:	FLOWCONTROL = 1
Default:	3
Normal Behavior:	0 = Disabled (No flow control capability) 1 = Receive Pause Frames (can receive and respond to PAUSE frames) 2 = Transmit Pause Frames (can send PAUSE frames) 3 = Both Enabled (can send and receive PAUSE frames) 0xFF = Hardware Default.
Possible Errors:	An error message is displayed if the value given is out of range.

USELASTSLOT

This parameter causes the driver to load on the device in the last slot found in the slot scan. The default behavior of the driver is to load on the first adapter found in the slot scan. This parameter forces the driver to load on the last one found instead.

Syntax:	UseLastSlot = [0 | any other value]
Example:	USELASTSLOT = 1
Default:	0
Normal Behavior:	0 = Disabled, any other value = Enabled
Possible Errors:	None

TXLOOPCOUNT

This parameter controls the number of times the transmit routine loops while waiting for a free transmit buffer. This parameter can affect Transmit performance.

Syntax:	TXLOOPCOUNT = <32-bit value>
Example:	TXLOOPCOUNT = 10000
Default:	1000
Normal Behavior:	Default
Possible Errors:	None

Data Center Bridging (DCB) for Intel® Network Connections

Data Center Bridging provides a lossless data center transport layer for using LANs and SANs in a single unified fabric.

Data Center Bridging includes the following capabilities:

- Priority-based flow control (PFC; IEEE 802.1Qbb)
- Enhanced transmission selection (ETS; IEEE 802.1Qaz)
- Congestion notification (CN)
- Extensions to the Link Layer Discovery Protocol standard (IEEE 802.1AB) that enable Data Center Bridging Capability Exchange Protocol (DCBX)

There are two supported versions of DCBX.

CEE Version: The specification can be found as a link within the following document: http://www.ieee802.org/1/files/public/docs2008/dcb-baseline-contributions-1108-v1.01.pdf

IEEE Version: The specification can be found as a link within the following document: https://standards.ieee.org/findstds/standard/802.1Qaz-2011.html

NOTE: The OS DCBX stack defaults to the CEE version of DCBX, and if a peer is transmitting IEEE TLVs, it will automatically transition to the IEEE version.

For more information on DCB, including the DCB Capability Exchange Protocol Specification, go to http://www.ieee802.org/1/pages/dcbridges.html

DCB for Windows Configuration:

Intel Ethernet Adapter DCB functions can be configured using Windows Device Manager. Open the adapter's property sheet and select the **Data Center** tab.

You can use the Intel® PROSet to perform the following tasks:

- **Display Status**:
  - Enhanced Transmission Selection
  - Priority Flow Control
  - FCoE Priority

  Non-operational status: If the Status indicator shows that DCB is non-operational, there may be a number of possible reasons:

  - DCB is not enabled - select the checkbox to enable DCB.
  - One or more of the DCB features is in a non-operational state. The features which contribute to the non-operational status are PFC and APP:FCoE.

  A non-operational status is most likely to occur when **Use Switch Settings** is selected or **Using Advanced Settings** is active. This is generally a result of one or more of the DCB features not getting successfully exchanged with the switch. Possible problems include:

- One of the features is not supported by the switch.
- The switch is not advertising the feature.
- The switch or host has disabled the feature (this would be an advanced setting for the host).

- Disable/enable DCB
- Troubleshooting information

Hyper-V (DCB and VMQ)

NOTE: Configuring a device in the VMQ + DCB mode reduces the number of VMQs available for guest OSes.

DCB for Linux

DCB is supported on RHEL6 or later or SLES11 SP1 or later. See your operating system documentation for specifics.

iSCSI Over DCB

Intel® Ethernet adapters support iSCSI software initiators that are native to the underlying operating system. Data Center Bridging is most often configured at the switch. If the switch is not DCB capable, the DCB handshake will fail but the iSCSI connection will not be lost.

NOTE: DCB does not install in a VM. iSCSI over DCB is only supported in the base OS. An iscsi initiator running in a VM will not benefit from DCB ethernet enhancements.

Microsoft Windows Configuration

iSCSI installation includes the installation of the iSCSI DCB Agent (iscsidcb.exe) user mode service. The Microsoft iSCSI Software Initiator enables the connection of a Windows host to an external iSCSI storage array using an Intel Ethernet adapter. Please consult your operating system documentation for configuration details.

Enable DCB on the adapter by the following:

1. From **Windows Device Manager**, expand **Networking Adapters** and highlight the appropriate adapter (such as Intel® Ethernet Server Adapter X520). Right click on the **Intel adapter** and select **Properties**.
2. In the **Property Page**, select the **Data Center** Tab.

The **Data Center** Tab provides feedback as to the DCB state, operational or non- operational, as well as providing additional details should it be non-operational.

Using iSCSI over DCB with ANS Teaming

The Intel® iSCSI Agent is responsible for maintaining all packet filters for the purpose of priority tagging iSCSI traffic flowing over DCB-enabled adapters. The iSCSI Agent will create and maintain a traffic filter for an ANS Team if at least one member of the team has an "Operational" DCB status. However, if any adapter on the team does not have an "Operational" DCB status, the iSCSI Agent will log an error in the Windows Event Log

for that adapter. These error messages are to notify the administrator of configuration issues that need to be addressed, but do not affect the tagging or flow of iSCSI traffic for that team, unless it explicitly states that the TC Filter has been removed.

Linux Configuration

In the case of Open Source distributions, virtually all distributions include support for an Open iSCSI Software Initiator and Intel® Ethernet adapters will support them. Please consult your distribution documentation for additional configuration details on their particular Open iSCSI initiator.

Intel® 82599 and X540-based adapters support iSCSI within a Data Center Bridging cloud. Used in conjunction with switches and targets that support the iSCSI/DCB application TLV, this solution can provide guaranteed minimum bandwidth for iSCSI traffic between the host and target. This solution enables storage administrators to segment iSCSI traffic from LAN traffic, similar to how they can currently segment FCoE from LAN traffic. Previously, iSCSI traffic within a DCB supported environment was treated as LAN traffic by switch vendors. Please consult your switch and target vendors to ensure that they support the iSCSI/DCB application TLV.

Remote Boot

Remote Boot allows you to boot a system using only an Ethernet adapter. You connect to a server that contains an operating system image and use that to boot your local system.

Flash Images

"Flash" is a generic term for nonvolatile RAM (NVRAM), firmware, and option ROM (OROM). Depending on the device, it can be on the NIC or on the system board.

Updating the Flash in Microsoft Windows

Intel® PROSet for Windows* Device Manager can update the flash on an Intel Ethernet network adapter. However, if you need to enable or disable the Boot ROM use BootUtil.

Intel® PROSet for Windows Device Manager can only be used to program add-in Intel Ethernet network adapters. LOM (LAN On Motherboard) network connections cannot be programmed with the UEFI network driver option ROM.

Using Intel PROSet to flash the UEFI Network Driver Option ROM

Intel® PROSet for Windows Device Manager can install the UEFI network driver on an Intel network adapter's option ROM. The UEFI network driver will load automatically during system UEFI boot when installed in the option ROM. UEFI specific *.FLB images are included in the downloaded release media. The "Boot Options" tab in Intel® PROSet for Windows Device Manager will allow the UEFI *.FLB image to be installed on the network adapter.

Updating the Flash from Linux

The BootUtil command line utility can update the flash on an Intel Ethernet network adapter. Run BootUtil with the following command line options to update the flash on all supported Intel network adapters. For example, enter the following command line:

```
bootutil64e -up=efi -all
```

BootUtil can only be used to program add-in Intel network adapters. LOM (LAN On Motherboard) network connections cannot be programmed with the UEFI network driver option ROM.

See the bootutil.txt file for details on using BootUtil.

Installing the UEFI Network Driver Option ROM from the UEFI Shell

The BootUtil command line utility can install the UEFI network driver on an Intel network adapter's option ROM. The UEFI network driver will load automatically during system UEFI boot when installed into the option ROM. For example, run BootUtil with the following command line options to install the UEFI network driver on all supported Intel network adapters:

```
FS0:\>bootutil64e -up=efi -all
```

BootUtil can only be used to program add-in Intel Ethernet network adapters. LOM (LAN On Motherboard) network connections cannot be programmed with the UEFI network driver option ROM.

See the bootutil.txt file for details on using BootUtil.

Enable Remote Boot

If you have an Intel Desktop Adapter installed in your client computer, the flash ROM device is already available in your adapter, and no further installation steps are necessary. For Intel Server Adapters, the flash ROM can be enabled using the BootUtil utility. For example, from the command line type:

```
BOOTUTIL -E
BOOTUTIL -NIC=1 -FLASHENABLE
```

The first line will enumerate the ports available in your system. Choose a port. Then type the second line, selecting the port you wish to enable. For more details, see the bootutil.txt file.

UEFI Network Device Driver for Intel® Ethernet Network Connections

UEFI Network Stack

As of UEFI 2.1 there are two network stack configurations under UEFI. The most common configuration is the PXE based network stack. The alternate network stack provides IPv4 TCP, UDP, and MTFTP network protocol support. As of UEFI 2.1 the PXE and IP-based network stacks cannot be loaded or operate simultaneously. The following two sections describe each UEFI network stack configuration.

Reference implementations of the PXE and IP based network stack source code are available for download at www.tianocore.org.

Loading the UEFI Network Driver

The network driver can be loaded using the UEFI shell "load" command:

```
load e3040e2.efi
```

Configuring UEFI Network Stack for PXE

The PXE (Preboot eXecution Environment) based UEFI network stack provides support for UEFI network boot loaders downloaded from a WFM compliant PXE server. Services which can be enabled include Windows Deployment Services (WDS), Linux network installation (Elilo), and TFTP file transfers. To enable UEFI PXE services the following network protocol drivers must be loaded with: snp.efi, bc.efi, and pxedhcp4.efi. These drivers can be loaded from the UEFI "load" shell command, but are often included as part of the UEFI system firmware. The UEFI shell command "drivers" can be used to determine if the UEFI PXE drivers are included in the UEFI implementation. The drivers command will output a table listing drivers loaded in the system. The following entries must be present in order to network boot a UEFI system over PXE:

DRV	VERSION	TYPE	CFG	DIAG	#D	#C	DRIVER NAME	IMAGE NAME
F5	00000010	D	-	-	2	-	Simple Network Protocol Driver	SNP
F7	00000010	D	-	-	2	-	PXE Base Code Driver	BC
F9	00000010	D	-	-	2	-	PXE DHCPv4 Driver	PxeDhcp4
FA	03004000	B	X	X	2	2	Intel(R) Network Connection 3.0.00	/e3000e2.efi

A network boot option will appear in the boot options menu when the UEFI PXE network stack and Intel UEFI network driver have been loaded. Selecting this
boot option will initiate a PXE network boot.

Configuring UEFI Network Stack for TCP/UDP/MTFTP

An IP-based network stack is available to applications requiring IP-based network protocols such as TCP, UDP, or MTFTP. The following UEFI network drivers must be built into the UEFI platform implementation to enable this stack: SNP (Simple Network Protocol), MNP (Managed Network Protocol), ARP, DHCP4, IPv4, ip4config, TCPv4, UDPv4, and MTFTPv4. These drivers will show up in the UEFI "drivers" command output if they are included in the platform UEFI implementation:

DRV	VERSION	TYPE	CFG	DIAG	#D	#C	DRIVER NAME	IMAGE NAME
F5	00000010	D	-	-	2	-	IP4 CONFIG Network Service Driver	Ip4Config
F7	00000010	D	-	-	2	-	Simple Network Protocol Driver	SNP
F8	00000010	D	-	-	2	-	ARP Network Service Driver	Arp
F9	00000010	D	-	-	2	-	Tcp Network Service Driver	Tcp4
FA	00000010	D	-	-	2	-	IP4 Network Service Driver	Ip4
FB	00000010	D	-	-	2	-	DHCP Protocol Driver	Dhcp4
FC	00000010	D	-	-	6	-	UDP Network Service Driver	Udp4
FD	00000010	D	-	-	2	-	MTFTP4 Network Service	Mtftp4
FE	00000010	B	-	-	2	6	MNP Network Service Driver	/mnp.efi
FF	03099900	B	X	X	2	2	Intel(R) Network Connection 3.0.00	/e3000e2.efi

The ifconfig UEFI shell command must be used to configure each network interface. Running "ifconfig -?" from the UEFI shell will display usage instructions for ifconfig.

Unloading the UEFI Network Driver

To unload a network driver from memory the UEFI "unload" command is used. The syntax for using the unload command is as follows: "unload [driver handle]", where driver handle is the number assigned to the driver in the far left column of the "drivers" output screen.

Force Speed and Duplex

The UEFI network driver supports forced speed and duplex capability. The force speed and duplex menu can be accessed with UEFI shell command "drvcfg":

```
drvcfg -s [driver handle] [control handle]
```

The following speed and duplex configurations can be selected:

- Autonegotiate (recommended)
- 100 Mbps, full duplex
- 100 Mbps, half duplex
- 10 Mbps, full duplex
- 10 Mbps, half duplex

The speed and duplex setting selected must match the speed and duplex setting of the connecting network port. A speed and duplex mismatch between ports will result in dropped packets and poor network performance. It is recommended to set all ports on a network to autonegotiate. Connected ports must be set to autonegotiate in order to establish a 1 gigabit per second connection.

Fiber-optic and 10 gigabit ethernet adapters do not support forced speed and duplex.

Diagnostic Capability

The UEFI network driver features built in hardware diagnostic tests. The diagnostic tests are called with the UEFI shell drvdiag command.

`drvdiag -s` -Performs a basic hardware register test.
`drvdiag -e` -Performs an internal loopback transmit and receive test.

UEFI Known Issues

Long Initialization Times

Long initialization times observed with Intel's UEFI driver are caused when the UNDI.Initialize command is called with the PXE_OPFLAGS_INITIALIZE_CABLE_DETECT flag set. In this case, UNDI.Initialize will try to detect the link state.

If the port is connected and link is up, initialize will generally finish in about 3.5 seconds (the time needed to establish link, dependent on link conditions, link speed and controller type) and returns PXE_STATFLAGS_COMMAND_COMPLETE. If the port is disconnected (link is down), initialize will complete in about 5 seconds and return PXE_STATFLAGS_INIIALIZED_NO_MEDIA (driver initializes hardware then waits for link and timeouts when link is not establish in 5 seconds).

When UNDI.Initialize is called with PXE_OPFLAGS_INITIALIZE_DO_NOT_DETECT_CABLE the function will not try to detect link status and will take less than 1 second to complete.

The behavior of UNDI.Initialize is described in UEFI specs 2.3.1: Initializing the network device will take up to four seconds for most network devices and in some extreme cases (usually poor cables) up to twenty seconds. Control will not be returned to the caller and the COMMAND_COMPLETE status flag will not be set until the adapter is ready to transmit.

Intel® Boot Agent Configuration

Boot Agent Client Configuration

The Intel® Boot Agent software provides configuration options that allow you to customize the behavior of the Intel Boot Agent software. You can configure the Intel Boot Agent in any of the following environments:

- A Microsoft* Windows* Environment
- A Microsoft* MS-DOS* environment
- A pre-boot environment (before operating system is loaded)

The Intel Boot Agent supports PXE in pre-boot, Microsoft Windows*, and DOS environments. In each of these environments, a single user interface allows you to configure PXE protocols on Intel® Ethernet Adapters.

Configuring the Intel® Boot Agent in a Microsoft Windows Environment

If you use the Windows operating system on your client computer, you can use Intel® PROSet for Windows* Device Manager to configure and update the Intel Boot Agent software. Intel PROSet is available through the device manager. Intel PROSet provides a special tab, called the *Boot Options* tab, used for configuring and updating the Intel Boot Agent software.

To access the **Boot Options** tab:

1. Open Intel PROSet for Windows Device Manager by opening the **System** Control Panel. On the **Hardware** tab, click **Device Manager.**
2. Select the appropriate adapter and click the **Boot Options** tab. If the tab does not appear, update your network driver.
3. The **Boot Options** tab shows a list of current configuration parameters and their corresponding values. Corresponding configuration values appear for the selected setting in a drop-down box.
4. Select a setting you want to change from the **Settings** selection box.
5. Select a value for that setting from the **Value** drop-down list.
6. Repeat the preceding two steps to change any additional settings.
7. Once you have completed your changes, click **Apply Changes** to update the adapter with the new values.

Configuring the Intel® Boot Agent in an MS-DOS Environment

Intel provides a utility, Intel® Ethernet Flash Firmware Utility (BootUtil) for installing and configuring the Intel Boot Agent using the DOS environment. See bootutil.txt for complete information.

Configuring the Intel® Boot Agent in a Pre-Boot PXE Environment

NOTE: Intel Boot Agent may be disabled in the BIOS.

You can customize the behavior of the Intel Boot Agent software through a pre-boot (operating system independent) configuration setup program contained within the adapter's flash ROM. You can access this pre-boot configuration setup program each time the client computer cycles through the boot process.

When the boot process begins, the screen clears and the computer begins its Power On Self Test (POST) sequence. Shortly after completion of the POST, the Intel Boot Agent software stored in flash ROM executes. The Intel Boot Agent then displays an initialization message, similar to the one below, indicating that it is active:

```
Initializing Intel(R) Boot Agent Version X.X.XX
PXE 2.0 Build 083
```

NOTE: This display may be hidden by the manufacturer's splash screen. Consult your manufacturer's documentation for details.

The configuration setup menu shows a list of configuration settings on the left and their corresponding values on the right. Key descriptions near the bottom of the menu indicate how to change values for the configuration settings. For each selected setting, a brief "mini-Help" description of its function appears just above the key descriptions.

1. Highlight the setting you need to change by using the arrow keys.
2. Once you have accessed the setting you want to change, press the spacebar until the desired value appears.
3. Once you have completed your changes, press **F4** to update the adapter with the new values. Any changed configuration values are applied as the boot process resumes.

The table below provides a list of configuration settings, their possible values, and their detailed descriptions:

Configuration Setting	Possible Values	Description
Network Boot Protocol	PXE (Preboot eXecution Environment)	Select PXE for use with Network management programs, such as LANDesk* Management Suite. ***NOTE*:** Depending on the configuration of the Intel Boot Agent, this parameter may not be changeable.
Boot Order	Use BIOS Setup Boot Order Try network first, then local drives Try local drives first, then network Try network only Try local drives only	Sets the boot order in which devices are selected during boot up if the computer does not have its own control method. If your client computer's BIOS supports the BIOS Boot Specification (BBS), or allows PnP-compliant selection of the boot order in the BIOS setup program, then this setting will always be **Use BIOS Setup Boot Order** and cannot be changed. In this case, refer to the BIOS setup manual specific to your client computer to set up boot options. If your client computer does not have a BBS- or PnP-compliant BIOS, you can select any one of the other possible values listed for this setting except for **Use BIOS Setup Boot Order**.
Legacy OS Wakeup Support. (**For 82559-based adapters only**)	0 = Disabled (Default Value) 1 = Enabled	If set to 1, the Intel Boot Agent will enable PME in the adapter's PCI configuration space during initialization. This allows remote wakeup under legacy operating systems that don't normally support it. Note that enabling this makes the adapter technically non-compliant with the ACPI specification, which is why the default is disabled.

NOTE: If, during PXE boot, more than one adapter is installed in a computer and you want to boot from the boot ROM located on a specific adapter, you can do so by moving the adapter to the top of the BIOS Boot Order or by disabling the flash on the other adapters.

While the configuration setup menu is displayed, diagnostics information is also displayed in the lower half of

the screen. This information can be helpful during interaction with Intel Customer Support personnel or your IT team members. For more information about how to interpret the information displayed, refer to Diagnostics Information for Pre-boot PXE Environments.

Intel Boot Agent Target/Server Setup

Overview

For the Intel® Boot Agent software to perform its intended job, there must be a server set up on the same network as the client computer. That server must recognize and respond to the PXE or BOOTP boot protocols that are used by the Intel Boot Agent software.

NOTE: When the Intel Boot Agent software is installed as an upgrade for an earlier version boot ROM, the associated server-side software may not be compatible with the updated Intel Boot Agent. Contact your system administrator to determine if any server updates are necessary.

Linux* Server Setup

Consult your Linux* vendor for information about setting up the Linux Server.

Windows* Deployment Services

Nothing is needed beyond the standard driver files supplied on the media. Microsoft* owns the process and associated instructions for Windows Deployment Services. For more information on Windows Deployment Services perform a search of Microsoft articles at: http://technet.microsoft.com/en-us/library/default.aspx

Intel® Boot Agent Messages

Message Text	Cause
Invalid PMM function number.	PMM is not installed or is not working correctly. Try updating the BIOS.
PMM allocation error.	PMM could not or did not allocate the requested amount of memory for driver usage.
Option ROM initialization error. 64-bit PCI BAR addresses not supported, AX=	This may be caused by the system BIOS assigning a 64-bit BAR (Base Address Register) to the network port. Running the BootUtil utility with the -64d command line option may resolve this issue.
PXE-E00: This system does not have enough free conventional memory. The	System does not have enough free memory to run PXE image. The Intel Boot Agent was unable to find enough free base memory (below 640K) to install the PXE client software. The system cannot boot via PXE in its current configuration. The error returns control to the BIOS and the system does not attempt to remote boot. If this error persists, try updating your system's BIOS to the most-recent version. Contact your system administrator or your computer vendor's customer support to resolve

Intel Boot Agent cannot continue.	the problem.
PXE-E01: PCI Vendor and Device IDs do not match!	Image vendor and device ID do not match those located on the card. Make sure the correct flash image is installed on the adapter.
PXE-E04: Error reading PCI configuration space. The Intel Boot Agent cannot continue.	PCI configuration space could not be read. Machine is probably not PCI compliant. The Intel Boot Agent was unable to read one or more of the adapter's PCI configuration registers. The adapter may be mis-configured, or the wrong Intel Boot Agent image may be installed on the adapter. The Intel Boot Agent will return control to the BIOSand not attempt to remote boot. Try to update the flash image. If this does not solve the problem, contact your system administrator or Intel Customer Support.
PXE-E05: The LAN adapter's configuration is corrupted or has not been initialized. The Intel Boot Agent cannot continue.	The adapter's EEPROM is corrupted. The Intel Boot Agent determined that the adapter EEPROM checksum is incorrect. The agent will return control to the BIOS and not attempt to remote boot. Try to update the flash image. If this does not solve the problem, contact your system administrator or Intel Customer Support.
PXE-E06: Option ROM requires DDIM support.	The system BIOS does not support DDIM. The BIOS does not support the mapping of the PCI expansion ROMs into upper memory as required by the PCI specification. The Intel Boot Agent cannot function in this system. The Intel Boot Agent returns control to the BIOS and does not attempt to remote boot. You may be able to resolve the problem by updating the BIOS on your system. If updating your system's BIOS does not fix the problem, contact your system administrator or your computer vendor's customer support to resolve the problem.
PXE-E07: PCI BIOS calls not supported.	BIOS-level PCI services not available. Machine is probably not PCI compliant.
PXE-E09: Unexpected UNDI loader error. Status == xx	The UNDI loader returned an unknown error status. xx is the status returned.
PXE-E20: BIOS extended memory copy	BIOS could not move the image into extended memory.

error.	
PXE-E20: BIOS extended memory copy error. AH == xx	Error occurred while trying to copy the image into extended memory. xx is the BIOS failure code.
PXE-E51: No DHCP or BOOTP offers received.	The Intel Boot Agent did not receive any DHCP or BOOTP responses to its initial request. Please make sure that your DHCP server (and/or proxyDHCP server, if one is in use) is properly configured and has sufficient IP addresses available for lease. If you are using BOOTP instead, make sure that the BOOTP service is running and is properly configured.
PXE-E53: No boot filename received.	The Intel Boot Agent received a DHCP or BOOTP offer, but has not received a valid filename to download. If you are using PXE, please check your PXE and BINL configuration. If using BOOTP, be sure that the service is running and that the specific path and filename are correct.
PXE-E61: Media test failure.	The adapter does not detect link. Please make sure that the cable is good and is attached to a working hub or switch. The link light visible from the back of the adapter should be lit.
PXE-EC1: Base-code ROM ID structure was not found.	No base code could be located. An incorrect flash image is installed or the image has become corrupted. Try to update the flash image.
PXE-EC3: BC ROM ID structure is invalid.	Base code could not be installed. An incorrect flash image is installed or the image has become corrupted. Try to update the flash image.
PXE-EC4: UNDI ID structure was not found.	UNDI ROM ID structure signature is incorrect. An incorrect flash image is installed or the image has become corrupted. Try to update the flash image.
PXE-EC5: UNDI ROM ID structure is invalid.	The structure length is incorrect. An incorrect flash image is installed or the image has become corrupted. Try to update the flash image.
PXE-EC6: UNDI driver image is invalid.	The UNDI driver image signature was invalid. An incorrect flash image is installed or the image has become corrupted. Try to update the flash image.

PXE-EC8: !PXE structure was not found in UNDI driver code segment.	The Intel Boot Agent could not locate the needed !PXE structure resource. An incorrect flash image is installed or the image has become corrupted. Try to update the flash image. This may also be caused by the system BIOS assigning a 64-bit BAR (Base Address Register) to the network port. Running the BootUtil utility with the -64d command line option may resolve this issue.
PXE-EC9: PXENV + structure was not found in UNDI driver code segment.	The Intel Boot Agent could not locate the needed PXENV+ structure. An incorrect flash image is installed or the image has become corrupted. Try to update the flash image.
PXE-M0F: Exiting Intel Boot Agent.	Ending execution of the ROM image.
This option has been locked and cannot be changed.	You attempted to change a configuration setting that has been locked by your system administrator. This message can appear either from within Intel® PROSet's Boot Options tab when operating under Windows* or from the Configuration Setup Menu when operating in a stand-alone environment. If you think you should be able to change the configuration setting, consult your system administrator.
PXE-M0E: Retrying network boot; press ESC to cancel.	The Intel Boot Agent did not successfully complete a network boot due to a network error (such as not receiving a DHCP offer). The Intel Boot Agent will continue to attempt to boot from the network until successful or until canceled by the user. This feature is disabled by default. For information on how to enable this feature, contact Intel Customer Support.

Intel Boot Agent Troubleshooting Procedures

Common Issues

The following list of problems and associated solutions covers a representative set of problems that you might encounter while using the Intel Boot Agent.

After booting, my computer experiences problems

After the Intel® Boot Agent product has finished its sole task (remote booting), it no longer has any effect on the client computer operation. Thus, any issues that arise after the boot process is complete are most likely not related to the Intel Boot Agent product.

If you are having problems with the local (client) or network operating system, contact the operating system manufacturer for assistance. If you are having problems with some application program, contact the application manufacturer for assistance. If you are having problems with any of your computer's hardware or with the BIOS, contact your computer system manufacturer for assistance.

Cannot change boot order

If you are accustomed to redefining your computer's boot order using the motherboard BIOS setup program, the default settings of the Intel Boot Agent setup program can override that setup. To change the boot sequence, you must first override the Intel Boot Agent setup program defaults. A configuration setup menu appears allowing you to set configuration values for the Intel Boot Agent. To change your computer's boot order setting, see Configuring the Boot Agent in a Pre-boot PXE Environment.

My computer does not complete POST

If your computer fails to boot with an adapter installed, but *does* boot when you remove the adapter, try moving the adapter to another computer and using BootUtil to disable the Flash ROM.

If this does not work, the problem may be occurring before the Intel Boot Agent software even begins operating. In this case, there may be a BIOS problem with your computer. Contact your computer manufacturer's customer support group for help in correcting your problem.

There are configuration/operation problems with the boot process

If your PXE client receives a DHCP address, but then fails to boot, you know the PXE client is working correctly. Check your network or PXE server configuration to troubleshoot the problem. Contact Intel Customer Support if you need further assistance.

POST hang may occur if two or more ports on Quad Port Server Adapters are configured for PXE

If you have an Intel® Gigabit VT Quad Port Server Adapter, Intel® PRO/1000 PT Quad Port LP Server Adapter, or an Intel® PRO/1000 PF Quad Port Server Adapter with two or more ports configured for PXE, you may experience POST hangs on some server systems. If this occurs the suggested workaround is move the adapter to another system and disable PXE on all but one port of the Adapter. You may also be able to prevent this problem by disabling any on-board SCSI or SAS controllers in your system BIOS.

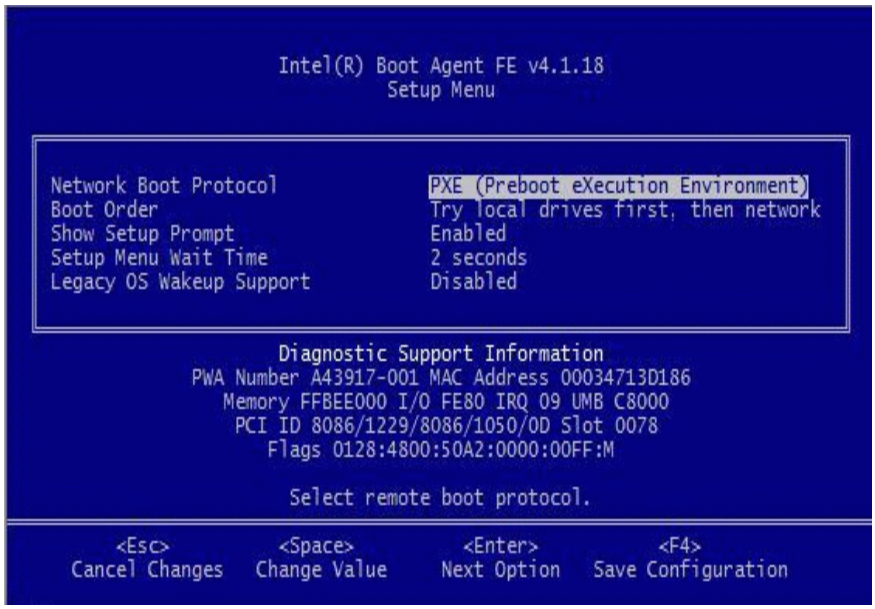
PXE option ROM does not follow the PXE specification with respect to the final "discover" cycle

In order to avoid long wait periods, the option ROM no longer includes the final 32-second discover cycle. (If there was no response in the prior 16-second cycle, it is almost certain that there will be none in the final 32-second cycle.

Diagnostics Information for Pre-boot PXE Environments

Anytime the configuration setup menu is displayed (see Configuring the Boot Agent in a Pre-boot PXE Environment), diagnostics information is also displayed on the lower portion of the screen. The information displayed appears similar to that shown in the lower half of the screen image below. This information can be helpful during interaction with Intel Customer Support personnel or your IT team members.

NOTE: Actual diagnostics information may vary, depending upon the adapter(s) installed in your computer.



Diagnostics information may include the following items:

Item	Description
PWA Number	The Printed Wire Assembly number identifies the adapter's model and version.
MAC Address	The unique Ethernet address assigned to the device.
Memory	The memory address assigned by the BIOS for memory-mapped adapter access.
I/O	The I/O port address assigned by the BIOS for I/O-mapped adapter access.
IRQ	The hardware interrupt assigned by the system BIOS.
UNB	The address in upper memory where the Boot Agent is installed by the BIOS.
PCI ID	The set of PCI identification values from the adapter in the form: VendorID/DeviceID/SubvendorID/SubdeviceID/Revision
Slot	The PCI bus address (slot number)reported by the BIOS. ***NOTE***: The number displayed is the BIOS version of the PCI slot number. Therefore, actual positions of NICs within physical slots may not be displayed as expected. Slots are not always enumerated in an obvious manner, and this will only report what is indicated by the BIOS.

Flags	A set of miscellaneous data either read from the adapter EEPROM or calculated by the Boot Agent initialization code. This information varies from one adapter to the next and is only intended for use by Intel customer support.

iSCSI Boot Configuration

iSCSI Initiator Setup

Configuring Intel® Ethernet iSCSI Boot on a Microsoft* Windows* Client Initiator

Requirements

1. Make sure the iSCSI initiator system starts the iSCSI Boot firmware. The firmware should be configured properly, be able to connect to iSCSI target, and detect the boot disk.
2. You will need Microsoft* iSCSI Software Initiator with integrated software boot support. This boot version of the initiator is available here.
3. To enable crash dump support, follow the steps in Crash Dump Support.

Configuring Intel® Ethernet iSCSI Boot on a Linux* Client Initiator

1. Install the Open-iSCSI initiator utilities.

   ```
   #yum -y install iscsi-inititator-utils
   ```

2. Refer to the README file found at https://github.com/mikechristie/open-iscsi.
3. Configure your iSCSI array to allow access.
   a. Examine */etc/iscsi/initiatorname.iscsi* for the Linux host initiator name.
   b. Update your volume manager with this host initiator name.
4. Set iscsi to start on boot.

   ```
   #chkconfig iscscd on
    #chkconfig iscsi on
   ```

5. Start iSCSI service (192.168.x.x is the IP Address of your target).

   ```
   #iscsiadm -n discovery -t s -p 192.168.x.x
   ```

   Observe the target names returned by iscsi discovery.
6. Log onto the target (-m XXX -T is XXX -l XXX -).

   ```
   iscsiadm -m node -T iqn.2123-01.com:yada:yada: -p 192.168.2.124
   -l
   ```

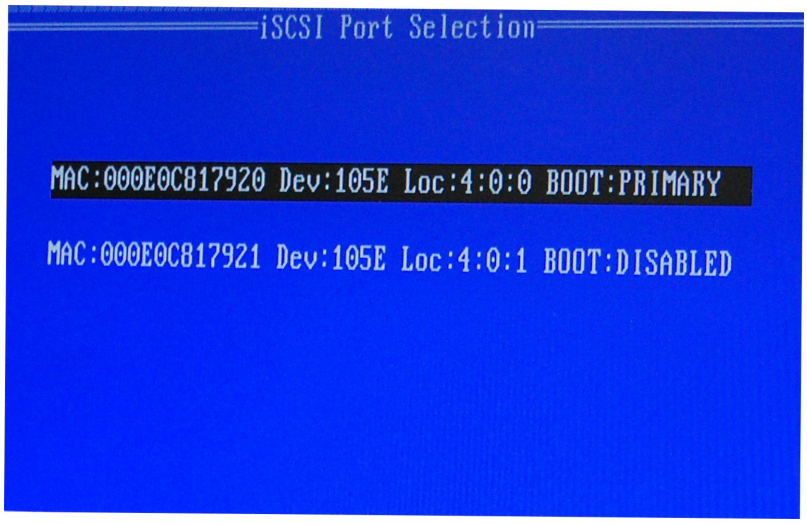
iSCSI Boot POST Setup

Intel® Ethernet iSCSI Boot features a setup menu which allows two network ports in one system to be enabled as iSCSI Boot devices. To configure Intel® iSCSI Boot, power-on or reset the system and press the Ctrl-D key when the message "`Press <Ctrl-D> to run setup...`" is displayed. After pressing the Ctrl-D key, you will be taken to the Intel® iSCSI Boot Port Selection Setup Menu.

NOTE: When booting an operating system from a local disk, Intel® Ethernet iSCSI Boot should be disabled for all network ports.

Intel® Ethernet iSCSI Boot Port Selection Menu

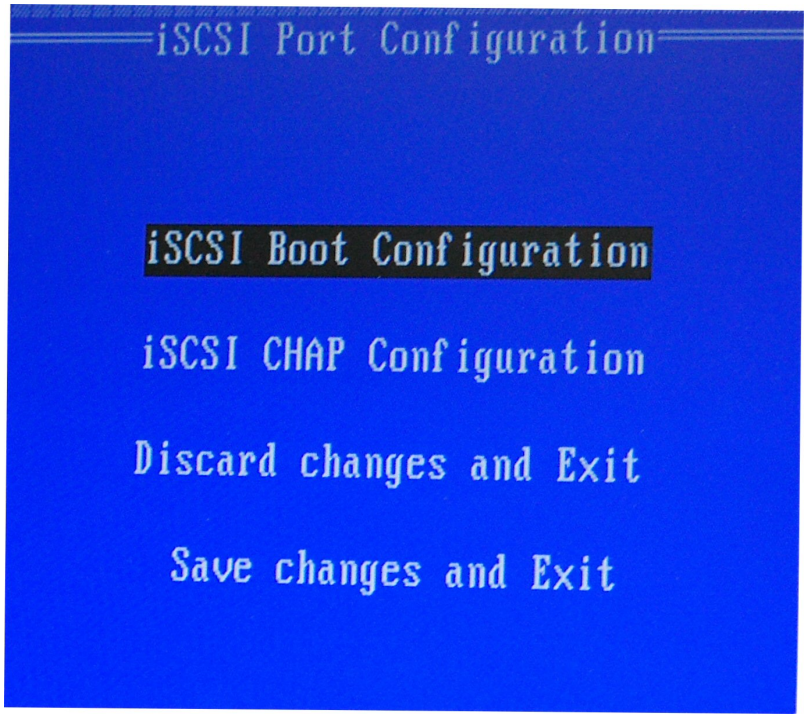
The first screen of the Intel® iSCSI Boot Setup Menu displays a list of Intel® iSCSI Boot-capable adapters. For each adapter port the associated PCI device ID, PCI bus/device/function location, and a field indicating Intel® Ethernet iSCSI Boot status is displayed. Up to 10 iSCSI Boot-capable ports are displayed within the Port Selection Menu. If there are more Intel® iSCSI Boot-capable adapters, these are not listed in the setup menu.



The usage of this menu is described below:

- One network port in the system can be selected as the primary boot port by pressing the 'P' key when highlighted. The primary boot port will be the first port used by Intel® Ethernet iSCSI Boot to connect to the iSCSI target. Only one port may be selected as a primary boot port.
- One network port in the system can be selected as the secondary boot port by pressing the 'S' key when highlighted. The secondary boot port will only be used to connect to the iSCSI target disk if the primary boot port fails to establish a connection. Only one port may be selected as a secondary boot port.
- Pressing the 'D' key with a network port highlighted will disable Intel® Ethernet iSCSI Boot on that port.
- Pressing the 'B' key with a network port highlighted will blink an LED on that port.
- Press the Esc key to leave the screen.

Intel® Ethernet iSCSI Boot Port Specific Setup Menu

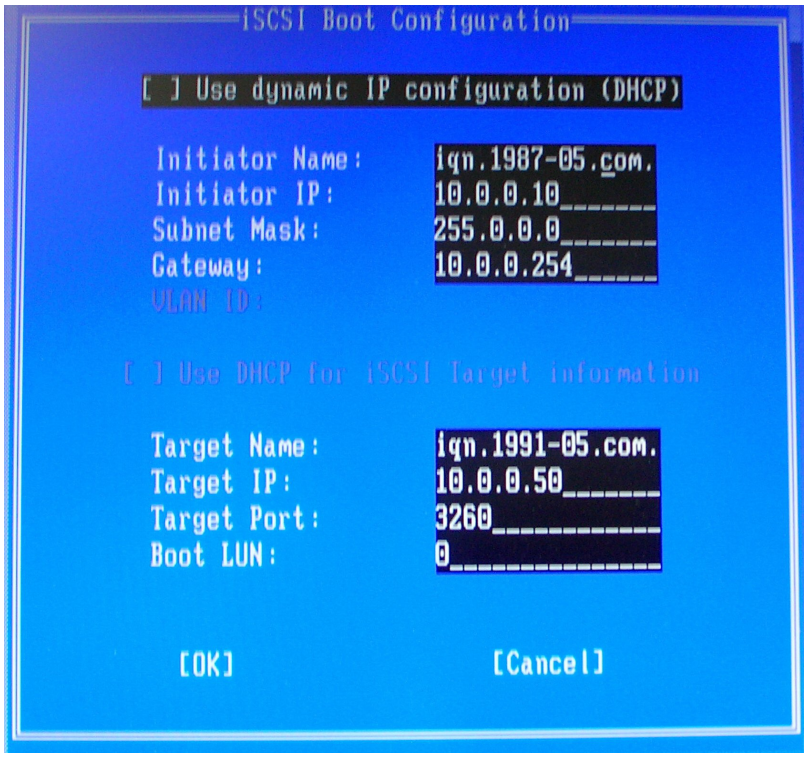


The port specific iSCSI setup menu has four options:

- **Intel® iSCSI Boot Configuration** - Selecting this option will take you to the iSCSI Boot Configuration Setup Menu. The iSCSI Boot Configuration Menu is described in detail in the section below and will allow you to configure the iSCSI parameters for the selected network port.
- **CHAP Configuration** - Selecting this option will take you to the CHAP configuration screen. The CHAP Configuration Menu is described in detail in the section below.
- **Discard Changes and Exit** - Selecting this option will discard all changes made in the iSCSI Boot Configuration and CHAP Configuration setup screens, and return back to the iSCSI Boot Port Selection Menu.
- **Save Changes and Exit** - Selecting this option will save all changes made in the iSCSI Boot Configuration and CHAP Configuration setup screens. After selecting this option, you will return to the iSCSI Boot Port Selection Menu.

Intel® iSCSI Boot Configuration Menu

The Intel® iSCSI Boot Configuration Menu allows you to configure the iSCSI Boot and Internet Protocol (IP) parameters for a specific port. The iSCSI settings can be configured manually or retrieved dynamically from a DHCP server.



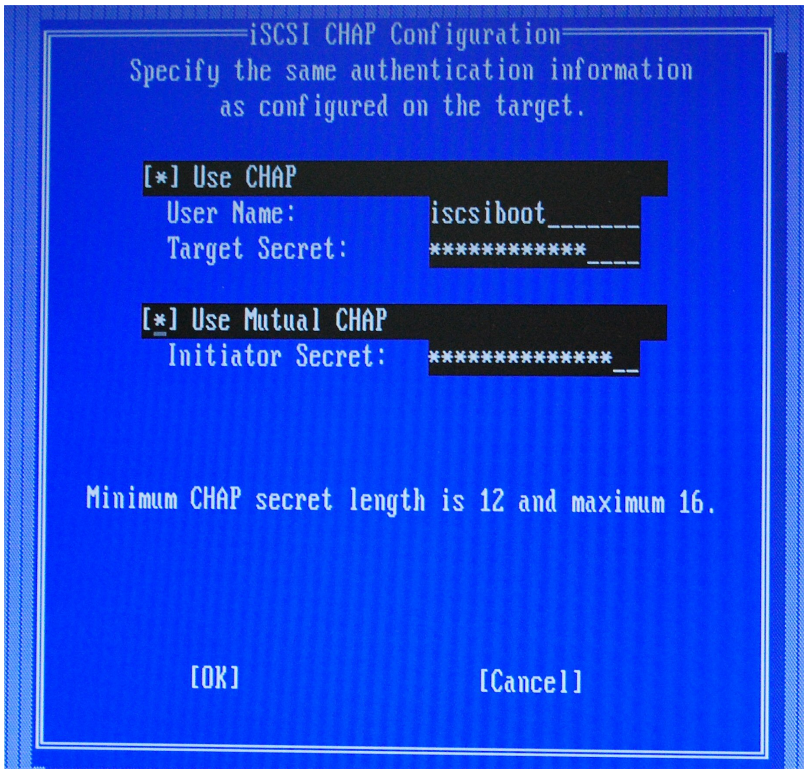
Listed below are the options in the Intel® iSCSI Boot Configuration Menu:

- **Use Dynamic IP Configuration (DHCP)** - Selecting this checkbox will cause iSCSI Boot to attempt to get the client IP address, subnet mask, and gateway IP address from a DHCP server. If this check-box is enabled, these fields will not be visible.
- **Initiator Name** - Enter the iSCSI initiator name to be used by Intel® iSCSI Boot when connecting to an iSCSI target. The value entered in this field is global and used by all iSCSI Boot-enabled ports in the system. This field may be left blank if the "`Use DHCP For Target Configuration`" checkbox is enabled. For information on how to retrieve the iSCSI initiator name dynamically from a DHCP server see the section DHCP Server Configuration.
- **Initiator IP** - Enter the client IP address to be used for this port as static IP configuration in this field. This IP address will be used by the port during the entire iSCSI session. This option is visible if DHCP is not enabled.
- **Subnet Mask** - Enter the IP subnet-mask in this field. This should be the IP subnet mask used on the network which the selected port will be connecting to for iSCSI. This option is visible if DHCP is not enabled.
- **Gateway IP** - Enter the IP address of the network gateway in this field. This field is necessary if the iSCSI target is located on a different sub-network than the selected Intel® iSCSI Boot port. This option is visible if DHCP is not enabled.
- **Use DHCP for iSCSI Target Information** - Selecting this checkbox will cause Intel® iSCSI Boot to attempt to gather the iSCSI target's IP address, IP port number, iSCSI target name, and SCSI LUN ID from a DHCP server on the network. For information on how to configure the iSCSI target parameters using DHCP see the section DHCP Server Configuration. When this checkbox is enabled, these fields will not be visible.
- **Target Name** - Enter the IQN name of the iSCSI target in this field. This option is visible if DHCP for iSCSI target is not enabled.

- **Target IP** - Enter the target IP address of the iSCSI target in this field. This option is visible if DHCP for iSCSI target is not enabled.
- **Target Port** - TCP Port Number.
- **Boot LUN -** Enter the LUN ID of the boot disk on the iSCSI target in this field. This option is visible if DHCP for iSCSI target is not enabled.

iSCSI CHAP Configuration

Intel® iSCSI Boot supports Mutual CHAP MD5 authentication with an iSCSI target. Intel® iSCSI Boot uses the "MD5 Message Digest Algorithm" developed by RSA Data Security, Inc.



The iSCSI CHAP Configuration menu has the following options to enable CHAP authentication:

- **Use CHAP** - Selecting this checkbox will enable CHAP authentication for this port. CHAP allows the target to authenticate the initiator. After enabling CHAP authentication, a user name and target password must be entered.
- **User Name** - Enter the CHAP user name in this field. This must be the same as the CHAP user name configured on the iSCSI target.
- **Target Secret** - Enter the CHAP password in this field. This must be the same as the CHAP password configured on the iSCSI target and must be between 12 and 16 characters in length. This password can not be the same as the **Initiator Secret.**
- **Use Mutual CHAP** – Selecting this checkbox will enable Mutual CHAP authentication for this port. Mutual CHAP allows the initiator to authenticate the target. After enabling Mutual CHAP authentication, an initiator password must be entered. Mutual CHAP can only be selected if Use CHAP is selected.
- **Initiator Secret -** Enter the Mutual CHAP password in this field. This password must also be configured on the iSCSI target and must be between 12 and 16 characters in length. This password can not be the same as the **Target Secret.**

The CHAP Authentication feature of this product requires the following acknowledgments:

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (http://www.openssl.org/).

Intel® PROSet for Windows* Device Manager

Many of the functions of the Intel® iSCSI Boot Port Selection Setup Menu can also be configured or revised from Windows Device Manager. Open the adapter's property sheet and select the **Data Options** tab. You must install the latest Intel Ethernet Adapter drivers and software to access this.

iSCSI Boot Target Configuration

For specific information on configuring your iSCSI target system and disk volume, refer to instructions provided by your system or operating system vendor. Listed below are the basic steps necessary to setup Intel® Ethernet iSCSI Boot to work with most iSCSI target systems. The specific steps will vary from one vendor to another.

NOTE: To support iSCSI Boot, the target needs to support multiple sessions from the same initiator. Both the iSCSI Boot firmware initiator and the OS High Initiator need to establish an iSCSI session at the same time. Both these initiators use the same Initiator Name and IP Address to connect and access the OS disk but these two initiators will establish different iSCSI sessions. In order for the target to support iSCSI Boot, the target must be capable of supporting multiple sessions and client logins.

1. Configure a disk volume on your iSCSI target system. Note the LUN ID of this volume for use when configuring in Intel® Ethernet iSCSI Boot firmware setup.
2. Note the iSCSI Qualified Name (IQN) of the iSCSI target, which will likely look like:

   ```
   iqn.1986-03.com.intel:target1
   ```

   This value is used as the iSCSI target name when you configuring your initiator system's Intel® Ethernet iSCSI Boot firmware.
3. Configure the iSCSI target system to accept the iSCSI connection from the iSCSI initiator. This usually requires listing the initiator's IQN name or MAC address for permitting the initiator to access to the disk volume. See the Firmware Setup section for information on how to set the iSCSI initiator name.
4. One-way authentication protocol can optionally be enabled for secure communications. Challenge-Handshake Authentication Protocol (CHAP) is enabled by configuring username/password on iSCSI target system. For setting up CHAP on the iSCSI initiator, refer to the section Firmware Setup for information.

Booting from Targets Larger than 2TB

You can connect and boot from a target LUN that is larger than 2 Terabytes with the following restrictions:

- The block size on the target must be 512 bytes
- The following operating systems are supported:
  - VMware* ESX 5.0, or later
  - Red Hat* Enterprise Linux* 6.3, or later
  - SUSE* Enterprise Linux 11SP2, or later
  - Microsoft* Windows Server* 2012 R2, or later
- You may be able to access data only within the first 2 TB.

NOTE: The Crash Dump driver does not support target LUNs larger than 2TB.

DHCP Server Configuration

If you are using Dynamic Host Configuration Protocol (DHCP), the DHCP server needs to be configured to provide the iSCSI Boot configurations to the iSCSI initiator. You must set up the DHCP server to specify Root Path option 17 and Host Name option 12 to respond iSCSI target information back to the iSCSI initiator. DHCP option 3, Router List may be necessary, depending on the network configuration.

DHCP Root Path Option 17:

The iSCSI root path option configuration string uses the following format:

```
iscsi:<server name or IP address>:<-
protocol>:<port>:<LUN>:<targetname>
```

- **Server name:** DHCP server name or valid IPv4 address literal.
  Example: `192.168.0.20`.
- **Protocol:** Transportation protocol used by ISCSI. Default is tcp (6).
  No other protocols are currently supported.
- **Port:** Port number of the iSCSI. A default value of 3260 will be used if this field is left blank.
- **LUN:** LUN ID configured on iSCSI target system. Default is zero.
- **Target name:** iSCSI target name to uniquely identify an iSCSI target in IQN format.
  Example: `iqn.1986-03.com.intel:target1`

DHCP Host Name Option 12:

Configure option 12 with the hostname of the iSCSI initiator.

DHCP Option 3, Router List:

Configure option 3 with the gateway or Router IP address, if the iSCSI initiator and iSCSI target are on different subnets.

Creating a Bootable Image for an iSCSI Target

There are two ways to create a bootable image on an iSCSI target:

- Install directly to a hard drive in an iSCSI storage array (Remote Install).
- Install to a local disk drive and then transfer this disk drive or OS image to an iSCSI Target (Local Install).

Microsoft* Windows*

Microsoft* Windows Server* natively supports OS installation to an iSCSI target without a local disk and also natively supports OS iSCSI boot. See Microsoft's installation instructions and Windows Deployment Services documentation for details.

SUSE* Linux Enterprise Server

For the easiest experience installing Linux onto an iSCSI target, you should use SLES10 or greater. SLES10 provides native support for iSCSI Booting and installing. This means that there are no additional steps outside of the installer that are necessary to install to an iSCSI target using an Intel Ethernet Server Adapter. Please refer to the SLES10 documentation for instructions on how to install to an iSCSI LUN.

Red Hat Enterprise Linux

For the easiest experience installing Linux onto an iSCSI target, you should use RHEL 5.1 or greater. RHEL 5.1 provides native support for iSCSI Booting and installing. This means that there are no additional steps outside of the installer that are necessary to install to an iSCSI target using an Intel Ethernet Server Adapter. Please refer to the RHEL 5.1 documentation for instructions on how to install to an iSCSI LUN.

Microsoft Windows Server iSCSI Crash Dump Support

Crash dump file generation is supported for iSCSI-booted Windows Server x64 by the Intel iSCSI Crash Dump Driver. To ensure a full memory dump is created:

1. Set the page file size equal to or greater than the amount of RAM installed on your system is necessary for a full memory dump.
2. Ensure that the amount of free space on your hard disk is able to handle the amount of RAM installed on your system.

To setup crash dump support follow these steps:

1. Setup Windows iSCSI Boot.
2. If you have not already done so, install the latest Intel Ethernet Adapter drivers and Intel PROSet for Windows Device Manager.
3. Open Intel PROSet for Windows Device Manager and select the Boot Options Tab.
4. From Settings select iSCSI Boot Crash Dump and the Value Enabled and click OK.

iSCSI Troubleshooting

The table below lists problems that can possibly occur when using Intel® Ethernet iSCSI Boot. For each problem a possible cause and resolution are provided.

Problem	Resolution
Intel® Ethernet iSCSI Boot does not load on system startup and the sign-on banner is not displayed.	• While the system logon screen may display for a longer time during system startup, Intel Ethernet iSCSI Boot may not be displayed during POST. It may be necessary to disable a system BIOS feature in order to display messages from Intel iSCSI Remote Boot. From the system BIOS Menu, disable any quiet boot or quick boot options. Also disable any BIOS splash screens. These options may be suppressing output from Intel iSCSI Remote Boot.

	• Intel Ethernet iSCSI Remote Boot has not been installed on the adapter or the adapter's flash ROM is disabled. Update the network adapter using the latest version of BootUtil as described in the Flash Images section of this document. If BootUtil reports the flash ROM is disabled, use the "`BOOTUTIL -flashenable`" command to enable the flash ROM and update the adapter. • The system BIOS may be suppressing output from Intel Ethernet iSCSI Boot. • Sufficient system BIOS memory may not be available to load Intel Ethernet iSCSI Boot. Attempt to disable unused disk controllers and devices in the system BIOS setup menu. SCSI controllers, RAID controller, PXE enabled network connections, and shadowing of system BIOS all reduce the memory area available to Intel Ethernet iSCSI Boot. Disable these devices and reboot the system to see if Intel iSCSI Boot is able to initialize. If disabling the devices in the system BIOS menu does not resolve the problem then attempt to remove unused disk devices or disk controllers from the system. Some system manufacturers allow unused devices to be disabled by jumper settings.
After installing Intel Ethernet iSCSI Boot, the system will not boot to a local disk or network boot device. The system becomes unresponsive after Intel Ethernet iSCSI Boot displays the sign-on banner or after connecting to the iSCSI target.	• A critical system error has occurred during iSCSI Remote Boot initialization. Power on the system and press the 's' key or 'ESC' key before Intel iSCSI Remote Boot initializes. This will bypass the Intel Ethernet iSCSI Boot initialization process and allow the system to boot to a local drive. Use the BootUtil utility to update to the latest version of Intel Ethernet iSCSI Remote Boot. • Updating the system BIOS may also resolve the issue.
"`Intel® iSCSI Remote Boot`" does not show up as a boot device in the system BIOS boot device menu.	• The system BIOS may not support Intel Ethernet iSCSI Boot. Update the system BIOS with the most recent version available from the system vendor. • A conflict may exist with another installed device. Attempt to disable unused disk and network controllers. Some SCSI and RAID controllers are known to cause compatibility problems with Intel iSCSI Remote Boot.
Error message displayed: "Failed to detect link"	• Intel Ethernet iSCSI Boot was unable to detect link on the network port. Check the link detection light on the back of the network connection. The link light should illuminate green when link is established with the link partner. If the link light is illuminated but the error message still displays then attempt to run the Intel link and cable diagnostics tests using `DIAGS.EXE` for DOS or Intel PROSet

	for Windows.
Error message displayed: "DHCP Server not found!"	iSCSI was configured to retrieve an IP address from DHCP but no DHCP server responded to the DHCP discovery request. This issue can have multiple causes: • The DHCP server may have used up all available IP address reservations. • The client iSCSI system may require static IP address assignment on the connected network. • There may not be a DHCP server present on the network. • Spanning Tree Protocol (STP) on the network switch may be preventing the Intel iSCSI Remote Boot port from contacting the DHCP server. Refer to your network switch documentation on how to disable Spanning Tree Protocol.
Error message displayed: "PnP Check Structure is invalid!"	• Intel Ethernet iSCSI Boot was not able to detect a valid PnP PCI BIOS. If this message displays Intel Ethernet iSCSI Boot cannot run on the system in question. A fully PnP compliant PCI BIOS is required to run Intel iSCSI Remote Boot.
Error message displayed: "Invalid iSCSI connection information"	• The iSCSI configuration information received from DHCP or statically configured in the setup menu is incomplete and an attempt to login to the iSCSI target system could not be made. Verify that the iSCSI initiator name, iSCSI target name, target IP address, and target port number are configured properly in the iSCSI setup menu (for static configuration) or on the DHCP server (for dynamic BOOTP configuration).
Error message displayed: "Unsupported SCSI disk block size!"	• The iSCSI target system is configured to use a disk block size that is not supported by Intel Ethernet iSCSI Boot. Configure the iSCSI target system to use a disk block size of 512 bytes.
Error message displayed: "ERROR: Could not establish TCP/IP connection with iSCSI target system."	• Intel Ethernet iSCSI Boot was unable to establish a TCP/IP connection with the iSCSI target system. Verify that the initiator and target IP address, subnet mask, port and gateway settings are configured properly. Verify the settings on the DHCP server if applicable. Check that the iSCSI target system is connected to a network accessible to the Intel iSCSI Remote Boot initiator. Verify that the connection is not being blocked by a firewall.
Error message displayed: "ERROR: CHAP authentication with target failed."	• The CHAP user name or secret does not match the CHAP configuration on the iSCSI target system. Verify the CHAP configuration on the Intel iSCSI Remote Boot port matches the iSCSI target system CHAP configuration. Disable CHAP in the iSCSI Remote Boot setup menu if it is not enabled on the target.

Error message displayed: "ERROR: Login request rejected by iSCSI target system."	• A login request was sent to the iSCSI target system but the login request was rejected. Verify the iSCSI initiator name, target name, LUN number, and CHAP authentication settings match the settings on the iSCSI target system. Verify that the target is configured to allow the Intel iSCSI Remote Boot initiator access to a LUN.
When installing Linux to NetApp Filer, after a successful target disk discovery, error messages may be seen similar to those listed below. Iscsi-sfnet:hostx: Connect failed with rc -113: No route to host Iscsi-sfnet:hostx: establish_session failed. Could not connect to target	• If these error messages are seen, unused iscsi interfaces on NetApp Filer should be disabled. • Continuous=no should be added to the iscsi.conf file
Error message displayed. "ERROR: iSCSI target not found."	• A TCP/IP connection was successfully made to the target IP address, however an iSCSI target with the specified iSCSI target name could not be found on the target system. Verify that the configured iSCSI target name and initiator name match the settings on the iSCSI target.
Error message displayed. "ERROR: iSCSI target can not accept any more connections."	• The iSCSI target cannot accept any new connections. This error could be caused by a configured limit on the iSCSI target or a limitation of resources (no disks available).
Error message displayed. "ERROR: iSCSI target has reported an error."	• An error has occurred on the iSCSI target. Inspect the iSCSI target to determine the source of the error and ensure it is configured properly.
Error message displayed. ERROR: There is an IP address conflict with another system on the network.	• A system on the network was found using the same IP address as the iSCSI Option ROM client. • If using a static IP address assignment, attempt to change the IP address to something which is not being used by another client on the network. • If using an IP address assigned by a DHCP server, make sure there are no clients on the network which are using an IP address which conflicts with the IP address range used by the DHCP server.

iSCSI Known Issues

A device cannot be uninstalled if it is configured as an iSCSI primary or secondary port.

Disabling the iSCSI primary port also disables the secondary port. To boot from the secondary port, change it to be the primary port.

iSCSI Remote Boot: Connecting back-to-back to a target with a Broadcom LOM

Connecting an iSCSI boot host to a target through a Broadcom LOM may occasionally cause the connection to fail. Use a switch between the host and target to avoid this.

iSCSI Remote Boot Firmware may show 0.0.0.0 in DHCP server IP address field

In a Linux base DHCP server, the iSCSI Remote Boot firmware shows 0.0.0.0 in the DHCP server IP address field. The iSCSI Remote Boot firmware looks at the DHCP server IP address from the Next-Server field in the DHCP response packet. However, the Linux base DHCP server may not set the field by default. Add "Next-Server <IP Address>;" in dhcpd.conf to show the correct DHCP server IP address.

iSCSI traffic stops after disabling RSC

To prevent a lost connection, Receive Segment Coalescing (RSC) must be disabled prior to configuring a VLAN bound to a port that will be used for connecting to an iSCSI target. Workaround this issue by disabling Receive Segment Coalescing before setting up the VLAN. This will avoid this traffic stop.

Microsoft Windows iSCSI Boot Issues

During OS install, injecting drivers from a DUP can cause a blue screen

During installation of Microsoft Windows Server 2012 on an iSCSI LUN, if you inject drivers from a DUP during the installation, you may experience a blue screen. Please install the hotfix described in kb2782676 to resolve the issue.

Microsoft Initiator does not boot without link on boot port:

After setting up the system for Intel® Ethernet iSCSI Boot with two ports connected to a target and successfully booting the system, if you later try to boot the system with only the secondary boot port connected to the target, Microsoft Initiator will continuously reboot the system.

To work around this limitation follow these steps:

1. Using Registry Editor, expand the following registry key:

   ```
   \System\CurrentControlSet\Services\Tcpip\Parameters
   ```

2. Create a DWORD value called DisableDHCPMediaSense and set the value to 0.

Support for Platforms Booted by UEFI iSCSI Native Initiator

Starting with version 2.2.0.0, the iSCSI crash dump driver gained the ability to support platforms booted using the native UEFI iSCSI initiator over supported Intel Network Adapters. This support is available on Windows Server or newer and only on x64 architecture. Any hotfixes listed above must also be applied.

Since network adapters on UEFI platforms may not provide legacy iSCSI option ROM, the boot options tab in DMIX may not provide the setting to enable the iSCSI crash dump driver. If this is the case, the following registry entry has to be created:

```
HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-
BFC1-08002BE10318}\<InstanceID>\Parameters
DumpMiniport    REG_SZ    iscsdump.sys
```

Moving iSCSI adapter to a different slot:

In a Windows* installation, if you move the iSCSI adapter to a PCI slot other than the one that it occupied when the drivers and MS iSCSI Remote Boot Initiator were installed, then a System Error may occur during the middle of the Windows Splash Screen. This issue goes away if you return the adapter to its original PCI slot. We recommend not moving the adapter used for iSCSI boot installation. This is a known OS issue.

If you have to move the adapter to another slot, then perform the following:

1. Boot the operating system and remove the old adapter
2. Install a new adapter into another slot
3. Setup the new adapter for iSCSI Boot
4. Perform iSCSI boot to the OS via the original adapter
5. Make the new adapter iSCSI-bootable to the OS
6. Reboot
7. Move the old adapter into another slot
8. Repeat steps 2 - 5 for the old adapter you have just moved

Uninstalling Driver can cause blue screen

If the driver for the device in use for iSCSI Boot is uninstalled via Device Manager, Windows will blue screen on reboot and the OS will have to be re-installed. This is a known Windows issue.

Adapters flashed with iSCSI image are not removed from the Device Manager during uninstall

During uninstallation all other Intel Network Connection Software is removed, but drivers for iSCSI Boot adapters that have boot priority.

I/OAT Offload may stop with Intel® Ethernet iSCSI Boot or with Microsoft Initiator installed

A workaround for this issue is to change the following registry value to "0":

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IOATDMA\Start

Only change the registry value if iSCSI Boot is enabled and if you want I/OAT offloading. A blue screen will occur if this setting is changed to "0" when iSCSI Boot is not enabled. It must be set back to "3" if iSCSI Boot is disabled or a blue screen will occur on reboot.

NDIS Driver May Not Load During iSCSI Boot F6 Install With Intel® PRO/1000 PT Server Adapter

If you are using two Intel® PRO/1000 PT Server Adapters in two PCI Express x8 slots of a rack mounted Xeon system, Windows installation can be done only via a local HDD procedure.

Automatic creation of iSCSI traffic filters for DCB, using Virtual Adapters created by Hyper-V, is only supported on Microsoft* Windows Server* 2008 releases R2 and later.

The iSCSI for Data Center Bridging (DCB) feature uses Quality of Service (QOS) traffic filters to tag outgoing packets with a priority. The Intel iSCSI Agent dynamically creates these traffic filters as needed for Windows Server 2008 R2 and later.

Invalid CHAP Settings May Cause Windows® Server 2008 to Blue Screen

If an iSCSI Boot port CHAP user name and secret do not match the target CHAP user name and secret, Windows Server 2008 may blue screen or reboot during installation or boot. Ensure that all CHAP settings match those set on the target(s).

Microsoft* Windows Server* 2008 Installation When Performing a WDS Installation

If you perform a WDS installation and attempt to manually update drivers during the installation, the drivers load but the iSCSI Target LUN does not display in the installation location list. This is a known WDS limitation with no current fix. You must therefore either perform the installation from a DVD or USB media or inject the drivers on the WDS WinPE image.

Microsoft has published a knowledge base case explaining the limitation in loading drivers when installing with iSCSI Boot via a WDS server.
http://support.microsoft.com/kb/960924

With high iSCSI traffic on Microsoft* Windows 2003 Server* R2, link flaps can occur with 82598-based silicon

This issue is caused by the limited support for Large Send Offload (LSO) in this Operating System. Please note that if ISCSI traffic is required for Windows 2003 Server R2, LSO will be disabled.

F6 Driver Does Not Support Standby Mode.

If you are performing an F6 Windows without a Local Disk installation, do not use Standby Mode.

F6 installation may fail with some EMC targets

An F6 installation may fail during the reboot in step 10 of “Installing Windows 2003 without a Local Disk” because of a conflict between the Intel F6 driver, the Microsoft iSCSI Initiator and the following EMC target model firmware versions:

- AX4-5 arrays: 02.23.050.5.705 or higher.
- CX300, CX500, CX700, and CX-3 Series arrays: 03.26.020.5.021 or higher.
- CX-4 Series arrays: 04.28.000.5.701 or higher, including all 04.29.000.5.xxx revisions.

To avoid the failure, ensure that the secondary iSCSI port cannot reach the target during the reboot in step 10.

iSCSI Boot and Teaming in Windows

Teaming is not supported with iSCSI Boot. Creating a team using the primary and secondary iSCSI adapters and selecting that team during the Microsoft initiator installation may fail with constant reboots. Do not select a team for iSCSI Boot, even if it is available for selection during initiator installation.

For load balancing and failover support, you can use MSFT MPIO instead. Check the Microsoft Initiator User Guide on how to setup MPIO.

Setting LAA (Locally Administered Address) on an iSCSI Boot-Enabled Port Will Cause System Failure on Next Reboot

Do not set LAA on ports with iSCSI Boot enabled.

Intel® Ethernet iSCSI Boot version does not match between displayed versions on DMIX and the scrolling text during boot

If a device is not set to primary but is enumerated first, the BIOS will still use that device's version of iSCSI Boot. Therefore the user may end up using an earlier version of Intel® Ethernet iSCSI Boot than expected. The solution is that all devices in the system must have the same version of iSCSI Boot. To do this the user should go to the Boot Options Tab and update the devices' flash to the latest version.

IPv6 iSCSI login to Intel Corporation EqualLogic arrays using jumbo frames

To establish an iSCSI session using IPv6 and jumbo frames with Dell EqualLogic arrays, TCP/UDP checksum offloads on the Intel iSCSI adapter should be disabled.

Microsoft Windows iSCSI/DCB Known Issues

iSCSI over DCB using Microsoft* Windows Server* 2012

iSCSI over DCB (priority tagging) is not possible on the port on which VMSwitch is created. This is by design in Microsoft* Windows Server* 2012.

Automatic creation of iSCSI traffic filters for DCB is only supported on networks which make use of IPv4 addressing

The iSCSI for Data Center Bridging (DCB) feature uses Quality of Service (QOS) traffic filters to tag outgoing packets with a priority. The Intel iSCSI Agent dynamically creates these traffic filters as needed on networks using IPv4 addressing.

IPv6 iSCSI login to Dell EqualLogic arrays using jumbo frames

To establish an iSCSI session using IPv6 and jumbo frames with Dell EqualLogic arrays, TCP/UDP checksum offloads on the Intel iSCSI adapter should be disabled.

Linux Known Issues

Channel Bonding

Linux Channel Bonding has basic compatibility issues with iSCSI Boot and should not be used.

Authentications errors on EqualLogic target may show up in dmesg when running Red Hat* Enterprise Linux 4

These error messages do not indicate a block in login or booting and may safely be ignored.

LRO and iSCSI Incompatibility

LRO (Large Receive Offload) is incompatible with iSCSI target or initiator traffic. A panic may occur when iSCSI traffic is received through the ixgbe driver with LRO enabled. The driver should be built and installed with:

```
# make CFLAGS_EXTRA=-DIXGBE_NO_LRO install
```

FCoE Boot Configuration

FCoE Client Setup

Installing and Configuring Intel® Ethernet FCoE Boot on a Microsoft* Windows* Client



WARNINGS:

- **Do not update the base driver via the Windows Update method**
  Doing so may render the system inoperable, generating a blue screen. The FCoE Stack and base driver need to be matched. The FCoE stack may get out of sync with the base driver if the base driver is updated via Windows Update. Updating can only be done via the Intel® Network Connections Installer.
- **If you are running Microsoft* Windows Server* 2012 R2, you must install KB2883200.**
  Failure to do so may result in an Error 1719 and a blue screen.

New Installation on a Windows Server* system

From the Intel downloaded media: Click the **FCoE/DCB checkbox** to install Intel® Ethernet FCoE Protocol Driver and DCB. The MSI Installer installs all FCoE and DCB components including Base Driver.

Microsoft Hotfixes

The following Microsoft hotfixes have been found to be needed for specific use cases:

Windows 2008 R2

- KB983554 - High-performance storage devices fix
- KB2708811 - Data corruption occurs under random write stress

Multipath I/O (MPIO)

Windows 2008 R2

- KB979743 - MPIO - write errors
- KB981379 - MS DSM - target issues

Windows 2008 R2 SP1

- KB2406705

  Configuring MPIO Timers: <http://technet.microsoft.com/en-us/library/ee619749(WS.10).aspx> contains additional information about these registry settings.

  Set the PathRecoveryInterval value to 60

Intel® Ethernet FCoE Configuration Using Intel® PROSet for Windows* Device Manager

Many FCoE functions can also be configured or revised using Intel PROSet for Windows* Device Manager, accessed from the **FCoE Properties** button within the **Data Center** tab. You can use Intel PROSet to perform the following tasks:

- Configure FCoE initiator specific settings
- Go to the corresponding port driver
- Review FCoE initiator information
- Obtain general information
- Review statistics
- Obtain information about the initiator
- Obtain information about attached devices
- FIP discovered VLANs and status

In addition, you can find some FCoE RSS performance settings under the Performance Options of the Advanced tab of the Network Adapter device properties. For additional information see Receive Side Scaling.

NOTES:

- From the **Boot Options** tab, the user will see the **Flash Information** button. Clicking on the **Flash Information** button will open the **Flash Information** dialog. From the **Flash Information** dialog, clicking on the **Update Flash** button allows Intel® iSCSI Remote Boot, Intel® Boot Agent (IBA), Intel® Ethernet FCoE Boot, EFI, and CLP to be written. The update operation writes a new image to the adapter's Flash and modifies the EEPROM, which may temporarily disable the operation of the Windows* network device driver. You might need to reboot the computer following this operation.
- You cannot update the flash image of a LOM; this button will be disabled.

1. Create a disk target (LUN) on an available Fibre Channel target. Configure this LUN to be accessible to the WWPN address of the initiator of the host being booted.
2. Make sure the client system starts the Intel® Ethernet FCoE Boot firmware. The firmware should be configured properly, be able to connect to Fibre Channel target, and detect the boot disk.

Intel® PROSet for Windows* Device Manager

Many of the functions of the Intel® Ethernet FCoE Boot Port Selection Setup Menu can also be configured or revised using Intel® PROSet for Windows Device Manager.

- Intel® Ethernet FCoE Boot version is displayed on the **Boot Options** tab if the combo image supports FCoE Boot.
- Intel® Ethernet FCoE Boot is an **Active Image** option if FCoE Boot is supported by the combo image.
- The **Active Image** setting enables/disables Intel® Ethernet FCoE Boot in the EEPROM.
- Intel® Ethernet FCoE Boot settings are displayed if **FCoE Boot** is the active image.

Installing Windows Server from a Remote Disk ("Diskless Install")

After the Option ROM is installed, if you wish to install the Windows Server operating system directly to the FCoE disk, do the following:

1. Locate the FCoE drivers in \APPS\FCOEBOOT\Winx64\Setup. Extract all zipped files and copy to a CD/DVD or USB media.
2. Boot the install media.
3. Perform a Custom install and proceed to the "Where do you want to install Windows?" screen.
4. Use **Load Driver** to load the FCoE drivers. Browse to the location you chose previously and load the following two drivers in the specified order:
    1. Intel(R) Ethernet Setup Driver for FCoE.
    2. Intel(R) Ethernet Virtual Storage Miniport Driver for FCoE.

    Note: the FCoE drivers will block any other network traffic from the FCoE-supported ports until after Step 7 in this section. Do not attempt to install an NDIS miniport for any FCoE-supported ports prior to Step 7 in this section.
5. You should now see the FCoE disk or disks appear in the list of available install targets. All disks accessible by this initiator through the boot port should appear.
6. Select the FCoE disk configured for boot in the Option ROM and continue the install until Windows is installed and you are at the desktop.

7. Follow the instructions for a New Installation on a Windows Server* system. This will install the networking drivers and configure the FCoE drivers to work with the networking drivers. Note that you cannot deselect the FCoE feature. You will be prompted to reboot at the end of the installation process.
8. Windows may prompt you to reboot once again after it returns to the desktop.

Installing Windows Server with Local Disk

After the Option ROM is installed, if you wish to install Windows Server with local disk, do the following:

1. Follow the instructions for a New Installation on a Windows Server* system.
2. Verify that the FCoE Boot disk is available in the **Fabric View** tab of Intel® PROSet for Windows Device Manager, and verify that you are online using Windows Disk Manager.
3. Open a command prompt, run the fcoeprep.bat batch file. To find the batch file, navigate to the \APPS\FCOEBOOT\Winx64\Setup directory.
4. Shut Windows down and capture the OS image to a local disk partition.
5. Transfer the image from the local hard drive to the FCoE target. This may be done from within the local Windows installation.
1. For Windows 2008 R2 SP1 only: Run bcdboot.exe from the local Windows installation to make the FCoE disk bootable.
    - If a System Reserved partition exists on the FCoE disk, type: `bcdboot F:\Windows /s E:`
      where E: is the FCoE System Reserved partition and F: is the FCoE partition with the Windows directory.
    - If a System Reserved partition does not exist, type: `bcdboot E:\Windows /s E:`
      where E: is the FCoE partition with the Windows directory.
6. Shut down and remove the local disk.
7. Configure the system BIOS to boot from the FCoE disk and boot.

NOTE: See Microsoft's documentation for more detailed instructions.

Upgrading Windows Drivers on an Intel® Ethernet FCoE-Booted System

Upgrading an FCoE-booted system can only be done via the Intel® Network Connections Installer. A reboot is required to complete the upgrade. You cannot upgrade a port's Windows driver and software package if the port is in the path to the virtual memory paging file and is also part of a Microsoft LBFO Team. To complete the upgrade, remove the port from the LBFO team and restart the upgrade.

Validation and Storage Certification

The software components for Intel® Ethernet FCoE are comprised of two major components: the Intel® Ethernet base driver and the Intel® Ethernet FCoE Driver. They are developed and validated as an ordered pair. You are strongly encouraged to avoid scenarios, either through upgrades or Windows update, where the Intel® Ethernet driver version is not the version released with the corresponding Intel® Ethernet FCoE driver. For more information, visit the download center.

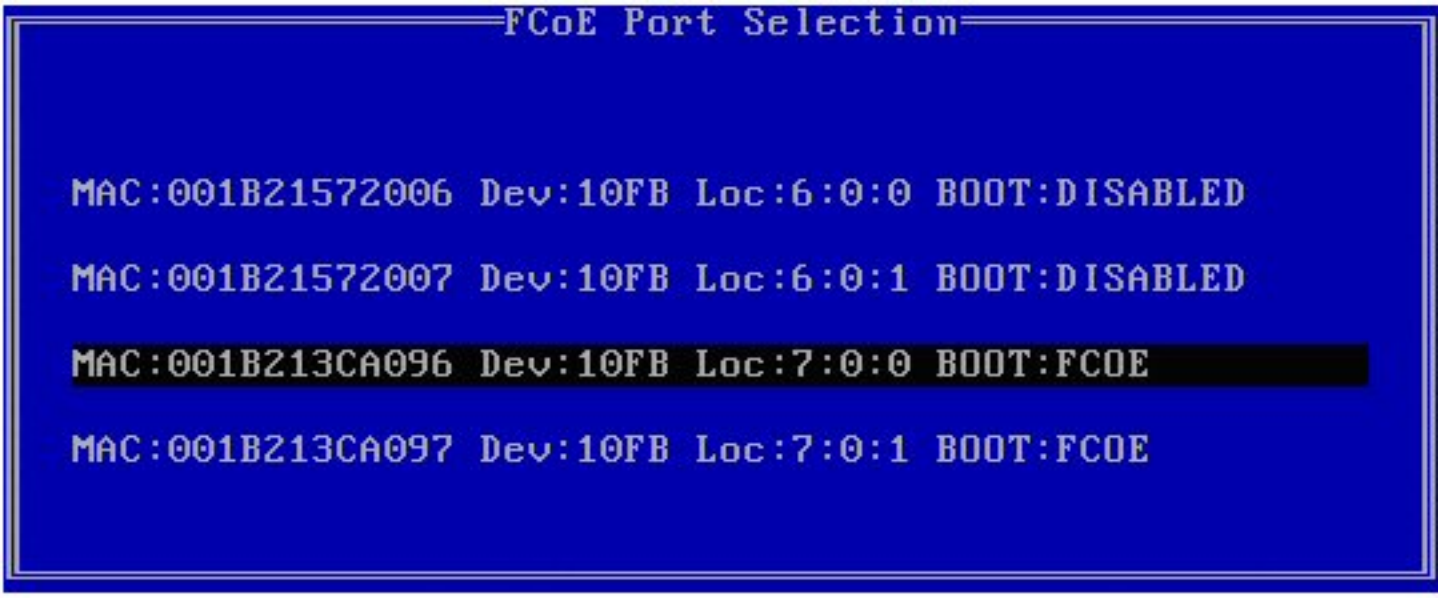
NOTES:

- Individually upgrading/downgrading the Intel® Ethernet FCoE driver will not work and may even cause a blue screen; the entire FCoE package must be the same version. Upgrade the entire FCoE package using the Intel® Network Connections installer only.
- If you uninstalled the Intel® Ethernet Virtual Storage Miniport Driver for FCoE component, just find the same version that you uninstalled and re-install it; or uninstall and then re-install the entire FCoE package.

Intel® Ethernet FCoE Boot Option ROM Setup

FCoE Port Selection Menu

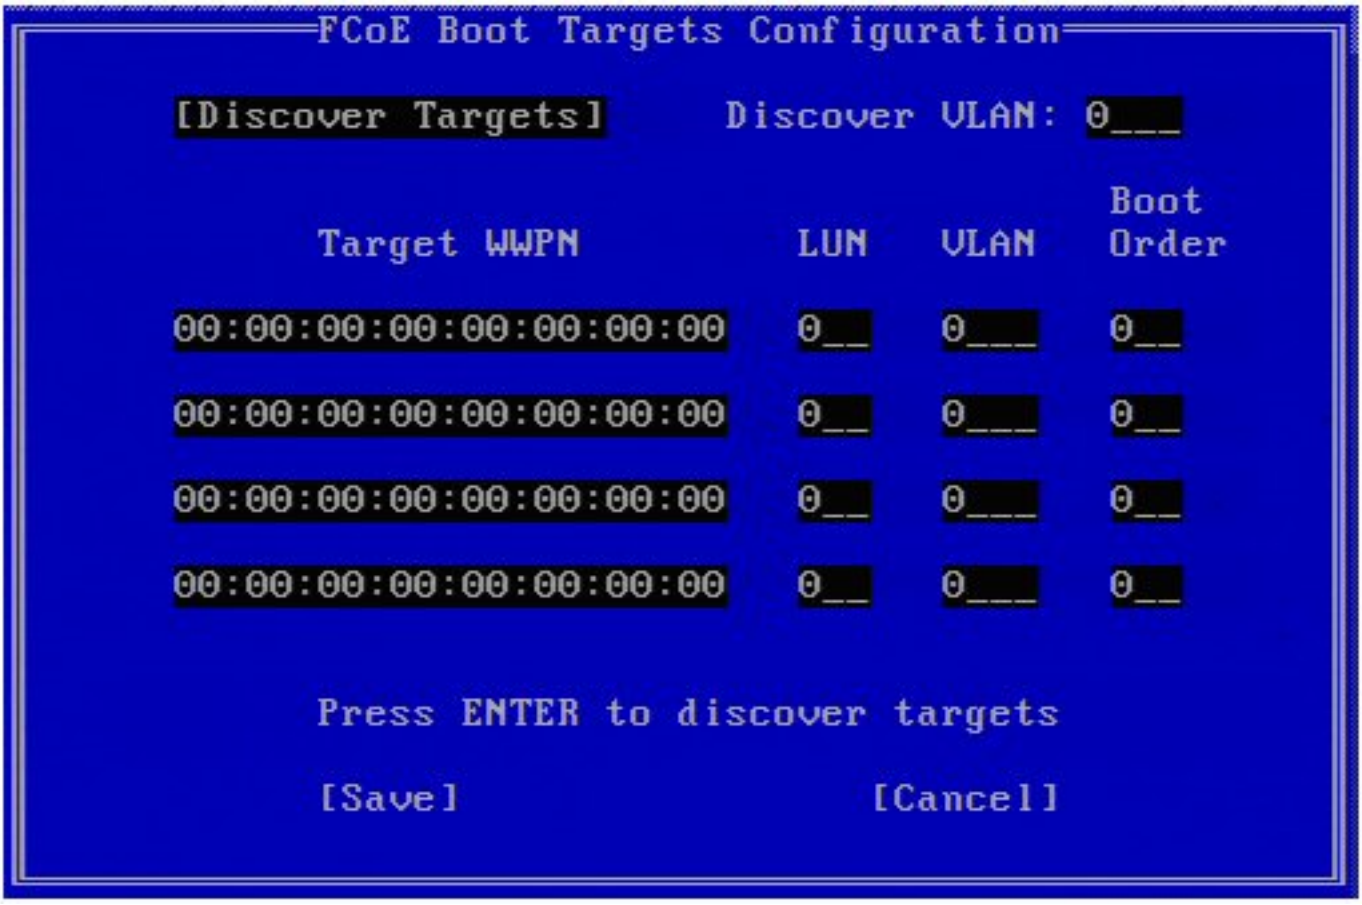
To configure Intel® Ethernet FCoE Boot, switch on or reset the system and press Ctrl-D when the message "`Press <Ctrl-D> to run setup...`" is displayed. After pressing Ctrl-D, you will be taken to the Intel® Ethernet FCoE Boot Port Selection Setup Menu.



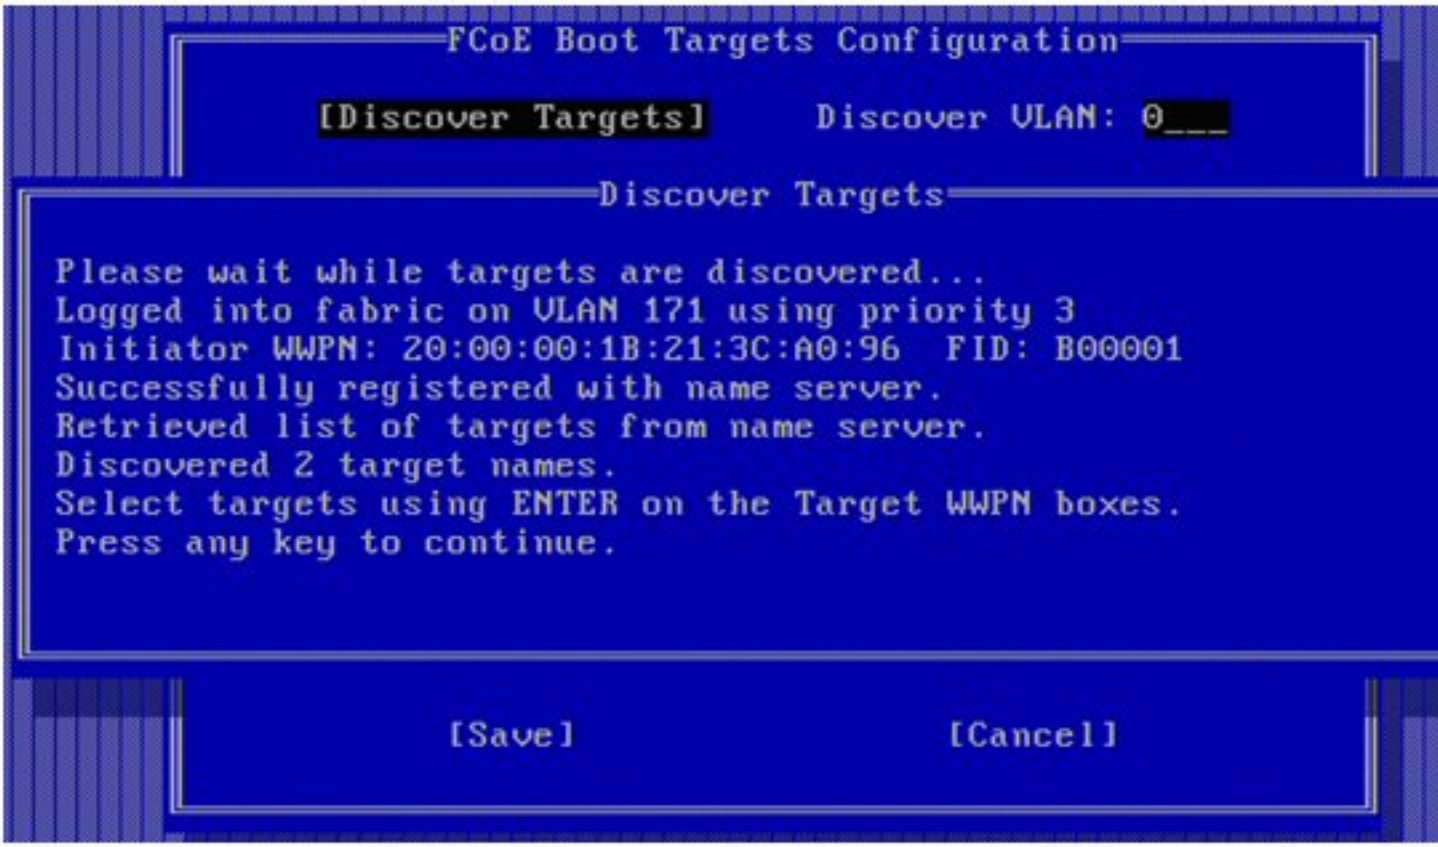
The first screen of the Intel® Ethernet FCoE Boot Setup Menu displays a list of Intel® FCoE Boot-capable adapters. For each adapter port, the associated SAN MAC address, PCI device ID, PCI bus/device/function location, and a field indicating FCoE Boot status is displayed. Up to 10 FCoE Boot-capable ports can be displayed within the Port Selection Menu. If there are more Intel® FCoE Boot-capable adapters, these are not listed in the setup menu.

Highlight the desired port and press **Enter**.

FCoE Boot Targets Configuration Menu



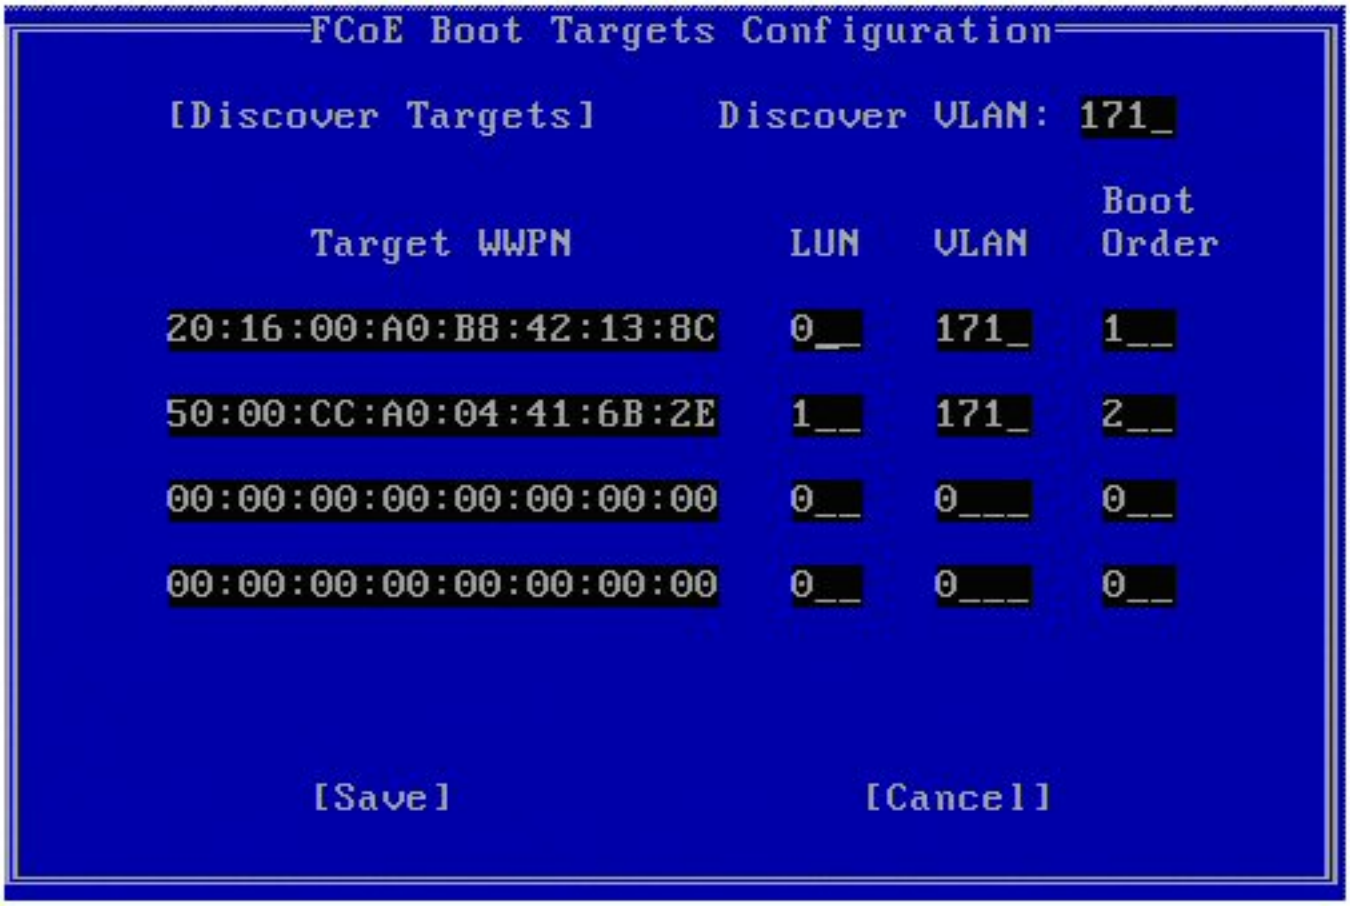
FCoE Boot Targets Configuration: Discover Targets is highlighted by default. If the **Discover VLAN** value displayed is not what you want, enter the correct value. Highlight **Discover Targets** and then press **Enter** to show targets associated with the **Discover VLAN** value. Under **Target WWPN**, if you know the desired WWPN you can manually enter it or press **Enter** to display a list of previously discovered targets.



FCoE Target Selection Menu



Highlight the desired **Target** from the list and press **Enter**.



Manually fill in the **LUN** and **Boot Order** values.

> **Boot Order** valid values are 0-4, where 0 means no boot order or ignore the target. A 0 value also indicates that this port should not be used to connect to the target. Boot order values of 1-4 can only be assigned once to target(s) across all FCoE boot-enabled ports.
>
> **VLAN** value is 0 by default. You may do a **Discover Targets** which will display a VLAN. If the VLAN displayed is not the one you require, enter the VLAN manually and then perform **Discover Targets** on that VLAN.

Hit **Save**.

NOTE: After the **Discover Targets** function is executed, the Option ROM will attempt to remain logged into the fabric until the **FCoE Boot Targets Configuration Menu** is exited.

- Keyboard Shortcuts: Up/Down, TAB and SHIFT-TAB to move between the controls. Left/Right/Home/End/Del/Backspace in the edit boxes.
- Press the Esc key to leave the screen.

UEFI Setup for Intel® FCoE Boot

Once you complete the configuration, it will be stored in the system's firmware memory.

Before beginning the configuration, update the adapter's UEFI FCoE Option ROM using the BootUtil tool and the latest BootIMG.FLB file. Use the following command:

```
BOOTUTIL64E.EFI -up=efi+efcoe -nic=PORT -quiet
```

where PORT is the NIC adapter number (in the following example nic=1)

NOTE: The UEFI FCoE driver must be loaded before you perform the following steps.

Accessing the FCoE Configuration Screen

Boot the system into its BIOS and proceed as follows:

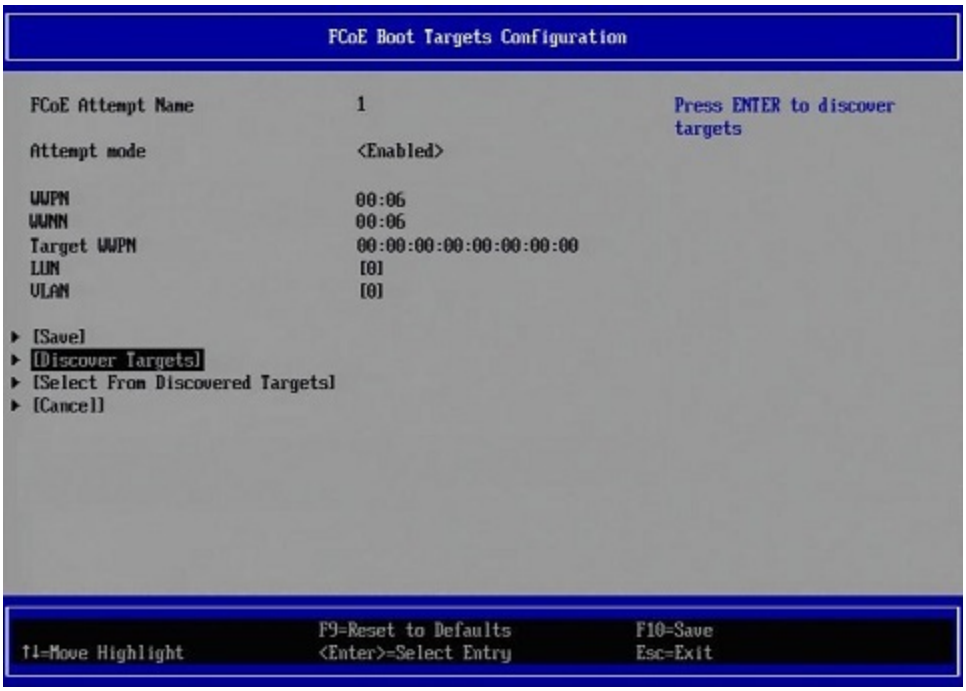
1. Select the **Advanced** tab, then choose **PCI Configuration**, then **UEFI Option ROM Control**, then **FCOE Configuration**.
2. The following screen is displayed:



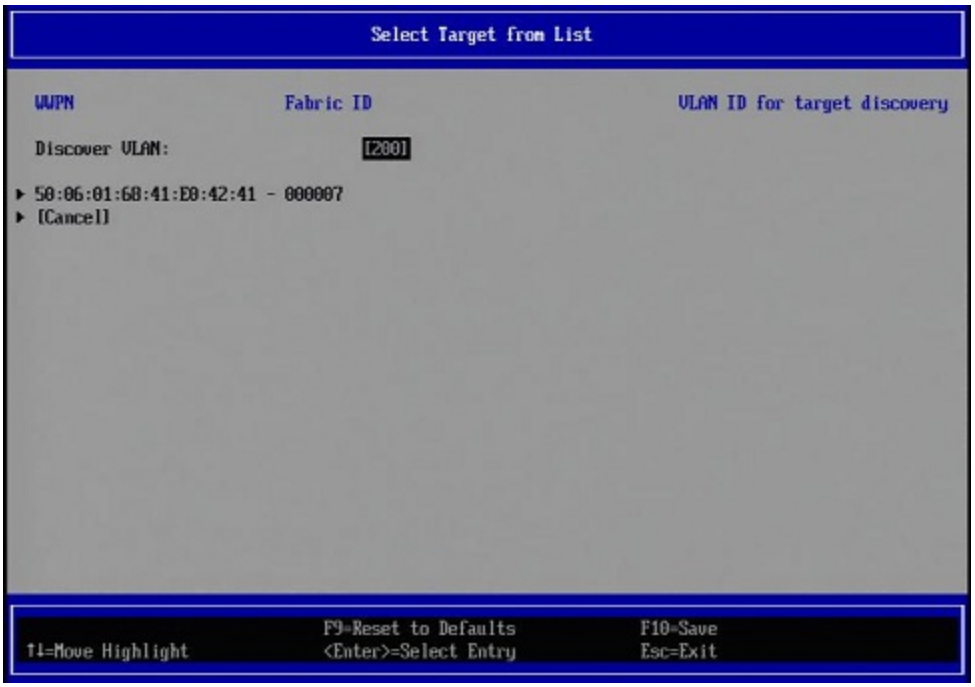
Adding an FCoE Attempt

An FCoE Attempt is a configured instance of a target from which the system will attempt to boot over FCoE.

1. From the FCoE Configuration menu, select **Add an Attempt**. All supported ports are displayed.
2. Select the desired port. The FCoE Boot Targets Configuration screen is displayed.



3. Select **Discover Targets** to automatically discover available targets (alternatively, you can manually enter the fields on the FCoE Boot Targets Configuration screen). The **Select from Discovered Targets** option displays a list of previously discovered targets.
4. Select **Auto-Discovery**. Note that the auto-discovery process may take several minutes. When auto-discovery is complete, the Select Target screen is displayed. **Discover VLAN** is the VLAN associated with a discovered adapter. There can be more than one target on a given VLAN.



5. Select the desired target from the list. The FCoE Boot Targets Configuration screen is displayed with completed fields for the selected target.
6. Press **F10 (Save)** to add this FCoE attempt. The FCoE Configuration screen is displayed with the newly added FCoE attempt listed.



Deleting an Existing FCoE Attempt

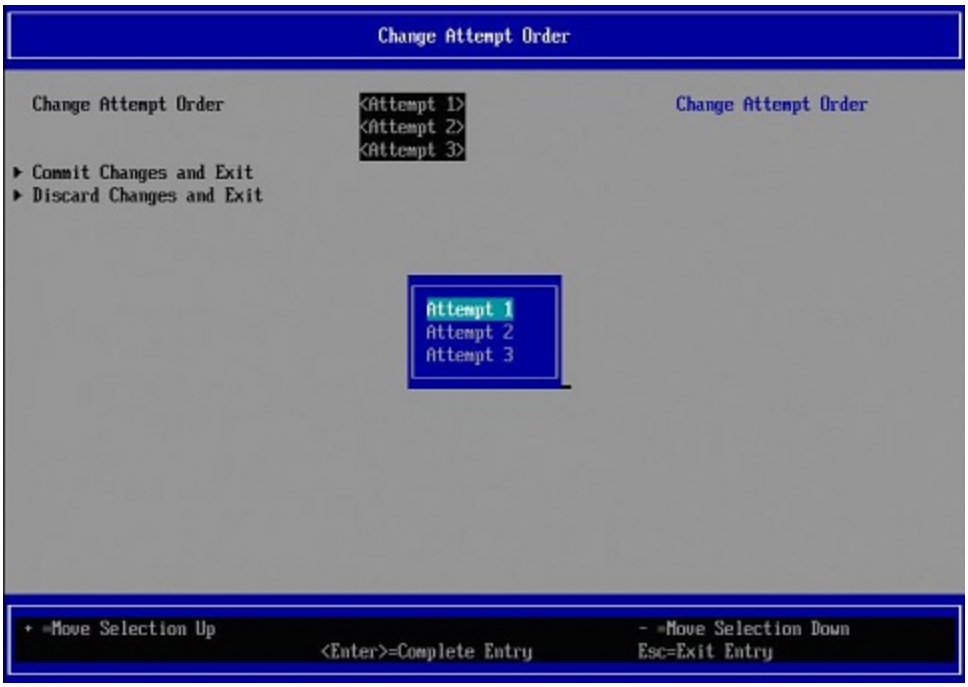
1. From the FCoE Configuration menu, select **Delete Attempts**.
2. Select one or more attempts to delete, as shown below (note that the example now shows three added attempts).



3. To delete the selected attempts, choose **Commit Changes and Exit**. To exit this screen without deleted the selected attempts, choose **Discard Changes and Exit**.

Changing the Order of FCoE Attempts

1. From the FCoE Configuration menu, select **Change Attempt Order**.
2. Press the Enter key to display the Change Attempt Order dialog, shown below.



3. Use the arrow keys to change the attempt order. When satisfied, press the Enter key to exit the dialog. The new attempt order is displayed.
4. To save the new attempt order, select **Commit Changes and Exit**. To exit without saving changes, select **Discard Changes and Exit**.

FCoE Boot Target Configuration

For specific information on configuring your FCoE target system and disk volume, refer to instructions provided by your system or operating system vendor. Listed below are the basic steps necessary to setup Intel® Ethernet FCoE Boot to work with most FCoE targets. The specific steps will vary from one vendor to another.

NOTE: If your FCoE Boot target is on a VLAN other than VLAN #1, then you must use the POST Boot Menu (Ctrl-D) to discover the target.

Installing Microsoft Windows Server from a Remote Disk ("Diskless Install")

After the Option ROM is installed, if you wish to install the Windows Server operating system directly to the FCoE disk, do the following:

1. Locate the FCoE drivers in \APPS\FCOEBOOT\Winx64\Setup. Extract all zipped files and copy to a CD/DVD or USB media.
2. Boot the install media.
3. Perform a Custom install and proceed to the "Where do you want to install Windows?" screen.

4. Use **Load Driver** to load the FCoE drivers. Browse to the location you chose previously and load the following two drivers in the specified order:
    1. Intel(R) Ethernet Setup Driver for FCoE.
    2. Intel(R) Ethernet Virtual Storage Miniport Driver for FCoE.

    Note: the FCoE drivers will block any other network traffic from the FCoE-supported ports until after Step 7. Do not attempt to install an NDIS miniport for any FCoE-supported ports until after Step 7.
5. You should now see the FCoE disk or disks appear in the list of available install targets. All disks accessible by this initiator through the boot port should appear.
6. Select the FCoE disk configured for boot in the Option ROM and continue the install until Windows is installed and you are at the desktop.
7. Follow the instructions for a new Windows Server installation. This will install the networking drivers and configure the FCoE drivers to work with the networking drivers. Note that you cannot deselect the FCoE feature. You will be prompted to reboot at the end of the installation process.
8. Windows may prompt you to reboot once again after it returns to the desktop.

Installing Windows Server with Local Disk

After the Option ROM is installed, if you wish to install Windows Server with local disk, do the following:

1. Follow the instructions for a new Windows Server installation.
2. Verify that the FCoE Boot disk is available in the **Fabric View** tab of Intel® PROSet for Windows Device Manager, and verify that you are online using Windows Disk Manager.
3. Open a command prompt, run the fcoeprep.bat batch file. To find the batch file, navigate to your architecture's directory within the \APPS\FCOEBOOT\Winx64\Setup directory.
4. Shut Windows down and capture the OS image to a local disk partition.
5. Transfer the image from the local hard drive to the FCoE target. This may be done from within the local Windows installation.
1. For Windows 2008 R2 SP1 only: Run bcdboot.exe from the local Windows installation to make the FCoE disk bootable.
    - If a System Reserved partition exists on the FCoE disk, type: `bcdboot F:\Windows /s E:`
      where E: is the FCoE System Reserved partition and F: is the FCoE partition with the Windows directory.
    - If a System Reserved partition does not exist, type: `bcdboot E:\Windows /s E:`
      where E: is the FCoE partition with the Windows directory.
6. Shut down and remove the local disk.
7. Configure the system BIOS to boot from the FCoE disk and boot.

NOTE: See Microsoft's documentation for more detailed instructions.

SUSE* Linux Enterprise Server

For the easiest experience installing Linux onto an FCoE target, you should use SLES11 or greater. SLES11 provides native support for FCoE Booting and installing. This means that there are no additional steps outside of the installer that are necessary to install to an iSCSI target using an Intel Ethernet Server Adapter. Please refer to the SLES11 documentation for instructions on how to install to an iSCSI LUN.

Red Hat Enterprise Linux

For the easiest experience installing Linux onto an iSCSI target, you should use RHEL 6 or greater. RHEL 6 provides native support for iSCSI Booting and installing. This means that there are no additional steps outside of the installer that are necessary to install to an iSCSI target using an Intel Ethernet Server Adapter. Please refer to the RHEL 6 documentation for instructions on how to install to an iSCSI LUN.

FCoE Known Issues

Intel® Ethernet FCoE Windows Issues

Intel® Ethernet Virtual Storage Miniport Driver for FCoE may disappear from Device Manager

The Intel® Ethernet Virtual Storage Miniport Driver for FCoE may disappear from the Device Manager after either:

- A virtual network is removed.
- The underlying Intel NIC adapter settings are modified.

This can occur when the corresponding Intel adapter is virtualized to create a new virtual network or delete or modify an existing Virtual Network. It can also happen when the underlying Intel NIC adapter settings are modified, including disabling or re-enabling the adapter.

As a workaround, remove all the resource dependencies of the Intel® Ethernet Virtual Storage Miniport Driver for FCoE that are currently being used by the system before making any changes to the Intel adapter for virtualization. For example, in one use case scenario, the user may have assigned the FCoE disk(s) from the FCoE storage driver to run one of its Virtual Machines, and at the same time the user wants to alter the configuration of the same Intel adapter for virtualization. In this scenario the user must remove the FCoE disks(s) from the Virtual Machine before altering the Intel adapter configuration.

Virtual Port may disappear from Virtual Machine

When the Virtual Machine starts, it asks the Intel® Ethernet Virtual Storage Miniport Driver for FCoE ("the driver") to create a Virtual Port. If the driver is subsequently disabled, the Virtual Port may disappear. The only way to get the Virtual Port back is to enable the driver and reboot the Virtual Machine.

When installing FCoE after installing ANS and creating AFT Team, Storports are not installed

If the user installs ANS and creates an AFT team and then installs FCoE/DCB, the result is that DCB is off by default. If the user then enables DCB on one port, the OS detects Storports and the user must manually click on the new hardware wizard prompts for each of them to install. If the user does not do that, DCB status is non-operational and the reason given is no peer.

Intel® PROSet for Windows Device Manager (DMiX) is not synched with FCoE CTRL-D Utility

When the user disables FCoE via the Control-D menu, the Intel PROSet for Windows Device Manager User Interface states that the flash contains an FCoE image, but that the flash needs to be updated. Updating the flash with the FCoE image again, re-enables FCoE and returns the user to the state where all the FCoE settings are available.

If the user uses the control-D menu to disable FCoE, then they should use the control-D menu to enable it because Intel PROSet for Windows Device Manager does not support enabling or disabling FCoE.

82599 and X540-based adapters don't display as SPC-3 compliant in Windows MPIO configuration

Because the FCoE initiator is a virtualized device it does not have its own unique hardware ID and thus is not displayed as a SPC-3 compliant device in Windows MPIO configuration.

When removing ALB teaming, all FCOE functions fail, all DMIX tabs are grayed out, and both adapter ports fail

For ANS teaming to work with Microsoft Network Load Balancer (NLB) in unicast mode, the team's LAA must be set to cluster node IP. For ALB mode, Receive Load Balancing must be disabled. For further configuration details, refer to http://support.microsoft.com/?id=278431

ANS teaming will work when NLB is in multicast mode, as well. For proper configuration of the adapter in this mode, refer to http://technet.microsoft.com/en-ca/library/cc726473(WS.10).aspx

FCoE and TCP/IP traffic on the same VLAN may not work on some switches

This is a known switch design and configuration issue.

Intel® Ethernet FCoE Boot Issues

Option ROM Known Issues

Discovery problems with multiple FCoE VLANs

The FCoE Option ROM may not discover the desired VLAN when performing VLAN discovery from the Discover Targets function. If the Discover VLAN box is populated with the wrong VLAN, then enter the desired VLAN before executing Discover Targets.

Windows Known Issues

Windows uses a paging file on the local disk

After imaging, if the local disk is not removed before booting from the FCoE disk then Windows may use the paging file from the local disk.

Crash dump to FCoE disks is only supported to the FCoE Boot LUN

The following scenarios are not supported:

- Crash dump to an FCoE disk if the Windows directory is not on the FCoE Boot LUN.
- Use of the DedicatedDumpFile registry value to direct crash dump to another FCoE LUN.

FCoE uninstall from a local disk may be blocked because installer inaccurately reports system is booted from FCoE

When the FCoE Option ROM connects to an FCoE disk during boot, the Windows installer may be unable to determine if the system was booted from FCoE or not and will block the FCoE uninstall. To uninstall, configure the Option ROM so that it does not connect to an FCoE disk.

Unable to create VLAN interfaces with Intel® Ethernet FCoE Boot enabled

When booted with FCoE, a user cannot create VLANs and/or Teams for other traffic types. This prevents converged functionality for non-FCoE traffic.

Server adapter configured for FCoE Boot available as External-Shared vnic via Hyper-V

If a port is set as a boot port, when the user installs the Hyper V role in the system and then goes into the Hyper V Network Manager to select which port to externally virtualize, the boot port displays, which it should not.

When setting the port to a boot port in Intel PROSet for Windows Device Manager, a message shows that the user should restart the system for the changes to be effective but does not force a restart. As a result the user level applications are in boot mode (i.e., Data Center Tab is grayed out) but kernel level drivers haven't been restarted to indicate to the OS that the port is a boot port. When the user then adds the Hyper V service to the system, the OS takes a snap shot of the ports available and this is the snap shot that it uses after the Hyper V role is added, system restarted and the user goes into the Hyper V Virtual Network Manager to virtualize the ports. As a result, the boot port also shows up.

Solutions:

Restart the system after setting a port to a boot port and before adding the Hyper V role. The port does not appear in the list of virtualizable ports in the Hyper V Virtual network manager.

Disable/enable the port in Device Manager after setting it to boot and before adding the Hyper V role. The port does not appear in the list of virtualizable ports in the Hyper V Virtual network manager.

FCoE Linkdown Timeout fails prematurely when Remote Booted

If an FCoE-booted port loses link for longer than the time specified in the **Linkdown Timeout** advanced setting in the Intel® Ethernet Virtual Storage Miniport Driver for FCoE, the system will crash. **Linkdown Timeout** values greater than 30 seconds may not provide extra time before a system crash.

Windows fails to boot properly after using the image install method

The following situation may arise when installing Windows for FCoE Boot using the imaging method: Windows boots successfully from the FCoE LUN when the local drive is installed, but when the local drive is removed, Windows seems to boot, but fails before reaching the desktop.

In this case it is likely that the Windows installation resides on both the FCoE LUN and local drive. This can be verified by booting from the FCoE LUN with the local drive installed, then comparing the drive letter in the path of files on the desktop with the drive letter for the boot partition in Windows' Disk Management tool. If the drive letters are different, then the Windows installation is split between the two disks.

If this situation is occurring, please ensure that **fcoeprep** is run prior to capturing the image, and that the system is not allowed to local boot between running **fcoeprep** and capturing the image. In addition, the local drive could be removed from the system prior to the first boot from the FCoE LUN.

Troubleshooting

Common Problems and Solutions

There are many simple, easy-to-fix problems related to network problems. Review each one of these before going further.

- Check for recent changes to hardware, software, or the network that may have disrupted communications.
- Check the driver software.
    - Make sure you are using the latest appropriate drivers for your adapter from the Intel support website.
    - Disable (or unload), then re-enable (reload) the driver or adapter.
    - Check for conflicting settings. Disable advanced settings such as teaming or VLANs to see if it corrects the problem.
    - Re-install the drivers.
- Check the cable. Use the best available cabling for the intended data rate.
    - Check that the cable is securely attached at both points.
    - Make sure the cable length does not exceed specifications.
    - 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.
    - Perform a cable test.
    - Replace the cable.
- Check the link partner (switch, hub, etc.).
    - Make sure the link partner is active and can send and receive traffic.
    - Make sure the adapter and link partner settings match one another, or are set to auto-negotiate.
    - Make sure the port is enabled.
    - Re-connect to another available port or another link partner.
- Look for adapter hardware problems.
    - Re-seat the adapter.
    - Insert the adapter in another slot.
    - Check for conflicting or incompatible hardware devices and settings.
    - Replace the adapter.
- Check the Intel support website for possible documented issues.
    - Select your adapter from the adapter family list.
    - Check the Frequently Asked questions section.
    - Check the Knowledge Base.
- Check your process monitor and other system monitors.
    - Check to see that there is sufficient processor and memory capacity to perform networking activity.
    - Look for any unusual activity (or lack of activity).
    - Use network testing programs to check for basic connectivity.

- Check your BIOS version and settings.
    - Use the latest appropriate BIOS for your computer.
    - Make sure the settings are appropriate for your computer.

The following troubleshooting table assumes that you have already reviewed the common problems and solutions.

Problem	Solution
Your computer cannot find the adapter	Make sure your adapter slots are compatible for the type of adapter you are using: • PCI Express v1.0 (or newer) • PCI-X v2.0 • PCI slots are v2.2
Diagnostics pass but the connection fails	Make sure the cable is securely attached, is the proper type and does not exceed the recommended lengths. 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 stops working after you installed the Intel® Network Adapter	Make sure your PCI BIOS is current. See PCI / PCI-X / PCI Express Configuration. 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 device drivers, then reload all drivers.
Adapter unable to connect to switch at correct speed. Gigabit adapter connects at 100 Mbps and 10 gigabit adapter connects at 1000 Mbps.	*This is applicable only to copper-based connections.* Make sure the adapter and the link partner are set to auto-negotiate. Verify that you are running the latest operating system revision for your switch and that the switch is compliant with the proper IEEE standard: • IEEE 802.3ad-compliant (gigabit over copper) • IEEE 802.3an-compliant (10 gigabit over copper)
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.
The adapter stops working without apparent cause	Run the adapter and network tests described under "Test the Adapter".
The Link indicator light is off	Run the adapter and network tests described under "Test the

Problem	Solution
	Adapter". Make sure the proper (and latest) driver is loaded. Make sure that the link partner is configured to auto-negotiate (or forced to match adapter) Verify that the switch is IEEE 802.3ad-compliant.
The link light is on, but communications are not properly established	Make sure the proper (and latest) driver is loaded. Both the adapter and its link partner must be set to either auto-detect 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.
RX or TX light is off	Network may be idle; try creating traffic while monitoring the lights.
The diagnostic utility reports the adapter is "Not enabled by BIOS"	The PCI BIOS isn't configuring the adapter correctly. See PCI / PCI-X / PCI Express Configuration.
The computer hangs when the drivers are loaded	Try changing the PCI BIOS interrupt settings. See PCI / PCI-X / PCI Express Configuration.
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.
PCI / PCI-X / PCI Express Configuration	If the adapter is not recognized by your OS or if it does not work you may need to change some BIOS settings. Try the following only if you are having problems with the adapter and are familiar with BIOS settings. • Check to see that the "Plug-and-Play" setting is compatible with the operating system you are using. • Make sure the slot is enabled. • Install the adapter in a bus-master slot. • Configure interrupts for level-triggering, as opposed to edge-triggering. • Reserve interrupts and/or memory addresses. This prevents multiple buses or bus slots from using the same interrupts. Check the BIOS for IRQ options for PCI /

Problem	Solution
	PCI-X / PCIe.

Multiple Adapters

When configuring a multi-adapter environment, you must upgrade all Intel adapters in the computer to the latest software.

If the computer has trouble detecting all adapters, consider the following:

- If you enable Wake on LAN* (WoL) on more than two adapters, the Wake on LAN feature may overdraw your system's auxiliary power supply, resulting in the inability to boot the system and other unpredictable problems. For multiple desktop/management adapters, it is recommended that you install one adapter at a time and use the IBAUtil utility (ibautil.exe in \APPS\BOOTAGNT) to disable the WoL feature on adapters that do not require WoL capabilities. On server adapters, the WoL feature is disabled by default.
- Adapters with 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).

Firmware Recovery Mode

When a device is in Firmware Recovery mode it will not pass traffic or allow any configuration; you can only attempt to recover the device's firmware. A device will enter Firmware Recovery mode if it detects a problem that requires the firmware to be reprogrammed.

Firmware Recovery Mode Errors and Messages

When a device is in Firmware Recovery mode, the device drivers and preboot software and utilities may log or display messages such as the following:

- Firmware recovery mode detected. Limiting functionality. Refer to the Intel(R) Ethernet Adapters and Devices User Guide for details on firmware recovery mode.
- Firmware recovery mode detected. Initialization failed.
- Firmware recovery mode detected. Limiting functionality.
- PXE-E0C: Firmware recovery mode detected. Initialization failed.

NOTE: To determine whether your device is in Firmware Recovery Mode, do the following:

- AC power cycle your system (see notes on AC power cycle in next section below).
- Once the system has booted to the operating system, wait approximately 10 minutes and then reboot the system.
- Once the system has booted to the operating system, the driver will indicate if the device is in Firmware Recovery Mode by logging error messages similar to the above. For Linux and ESX, check dmesg. For Microsoft Windows, check the Event Viewer.

Resolving Firmware Recovery Mode Issues

If your device is in Firmware Recovery mode you can restore it to factory defaults using the latest NVM Update Package. Download the latest NVM Update Package from your vendor's support website and follow

the instructions in it.

NOTES:

- Before initiating device recovery, the integrity of the host operating system, device drivers and firmware utilities must be verified and reinstalled if necessary. Fully functional operating system, device drivers and tools are required for device recovery.
- AC power cycle your system even if the instructions do not indicate that an AC power cycle is required. A power cycle is required to enable the device to restore and reload critical device settings to enable recovery. To perform an AC power cycle, do the following:
  - Shut down the system if it is powered up.
  - Unplug all AC power cords from the system.
  - Leave the AC power cords unplugged for 15 seconds to allow the system power supply to discharge completely.
  - Plug in AC power cords to the system.

Other Performance Issues

Attaining the best speed 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 often provide better results. Check for loose or damaged connectors. Check the cable for kinked or damaged sections.
- **Bus speed and traffic -** The PCI bus speed accommodates the slowest PCI card installed. Check to see if you have a card that is slowing down your system.
- **Processor and Memory** - Check your performance monitoring programs to see if traffic is being affected by your processor speed, available memory or other processes.
- **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 for your OS.
- **Operating System** - Networking feature implementation will vary by operating system version, such as offloading and multiprocessor threading.

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.

Testing from Windows

Intel PROSet allows you to run three 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: Determine 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 dependent on the adapter and operating system. Tests may be disabled if:

- iSCSI Boot is enabled on the port.
- FCoE Boot is enabled on the port.
- The port is used as a manageability port.
- The tests are being run from a virtual machine.

Testing from Windows PowerShell*

Intel provides two PowerShell cmdlets for testing your adapter.

- Test-IntelNetDiagnostics runs the specified test suite on the specified device. See the Test-IntelNetDiagnostics help inside PowerShell for more information.
- Test-IntelNetIdentifyAdapter blinks the LED on the specified device.

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.

DOS Diagnostics

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

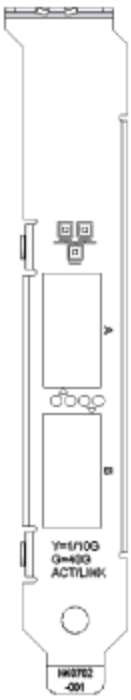
Spanning trees can be troublesome in a networking configuration. The Detect Spanning Tree option attempts to detect if a spanning tree exists on the network. This is done by resetting the link and listening for spanning tree packets.

Indicator Lights

The Intel Server and Desktop network adapters feature indicator lights on the adapter backplate that serve to indicate activity and the status of the adapter board. The following tables define the meaning for the possible states of the indicator lights for each adapter board.

Dual Port QSFP+ Adapters

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.

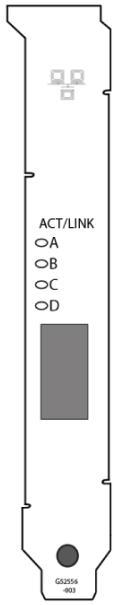
Single Port QSFP+ Adapters

The Intel® **Ethernet Converged Network Adapter XL710-Q1** 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 Converged Network Adapter X520-Q1** has the following indicator lights:



Label	Indication	Meaning
ACT/LNK	Green	Linked at 10 Gb
	Yellow	Linked at 1 Gb
	Blinking On/OFF	Actively transmitting or receiving data
	Off	No link.

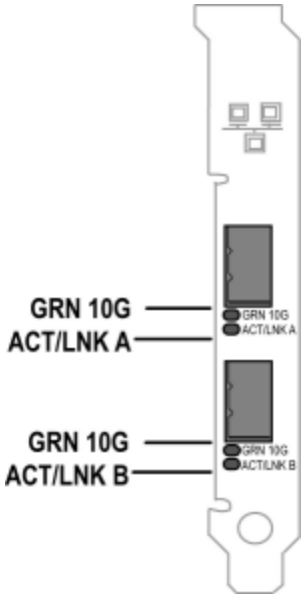
Dual Port SFP/SFP+ Adapters

The **Intel® Ethernet Converged Network Adapter X710-2** has the following indicator lights:



Label	Indication	Meaning
LNK	Green	Linked at 10 Gb
	Yellow	Linked at 1 Gb
ACT	Blinking On/OFF	Actively transmitting or receiving data
	Off	No link.

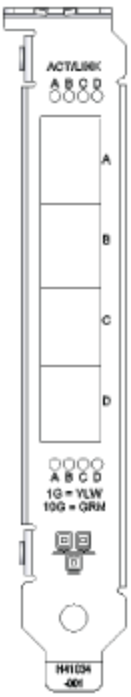
The **Intel® 10 Gigabit AF DA Dual Port Server Adapter** and **Intel® Ethernet Server Adapter X520** series of adapters have the following indicator lights:



Label	Indication	Meaning
GRN 10G (A or B): Green	On	Linked to the LAN.
	Off	Not linked to the LAN.
ACT/LNK (A or B): Green	Blinking On/Off	Actively transmitting or receiving data.
	Off	No link.

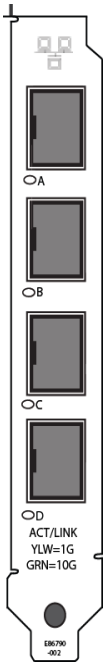
Quad Port SFP/SFP+ Adapters

The Intel® **Ethernet Converged Network Adapter X710-4** has the following indicator lights:



Label	Indication	Meaning
LNK	Green	Linked at 10 Gb
	Yellow	Linked at 1 Gb
ACT	Blinking On/OFF	Actively transmitting or receiving data
	Off	No link.

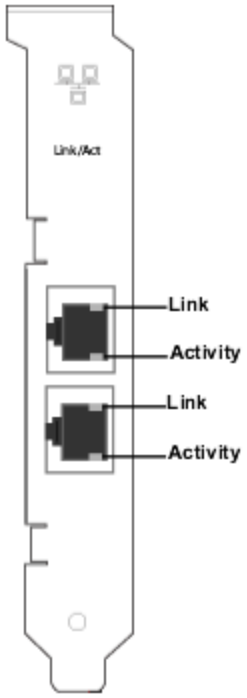
The Intel® **Ethernet Converged Network Adapter X520-4** has the following indicator lights:



Label	Indication	Meaning
ACT/L NK	Green	Linked at 10 Gb
	Yellow	Linked at 1 Gb
	Blinking On/OFF	Actively transmitting or receiving data
	Off	No link.

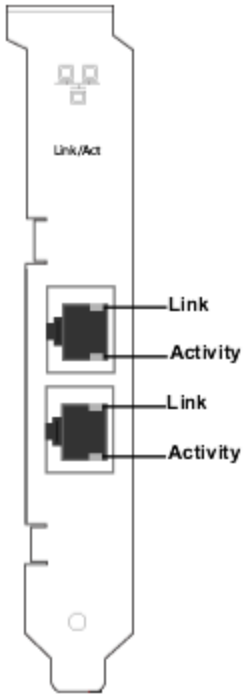
Dual Port Copper Adapters

The **Intel® Ethernet Converged Network Adapter X550-T2** has the following indicator lights:



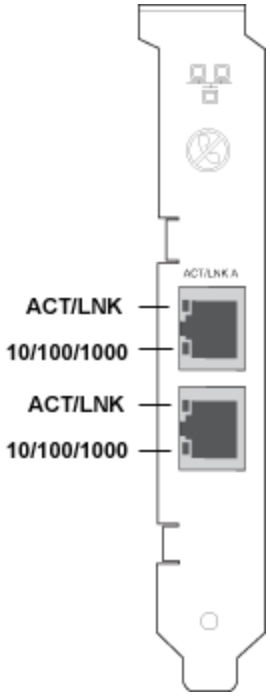
Label	Indication	Meaning
Link	Green	Linked at 10 Gb.
	Yellow	Linked at 1 Gb or 100 Mbps.
	Off	No link.
Activity	Blinking On/Off	Actively transmitting or receiving data.
	Off	No link.

The **Intel® Ethernet Server Adapter X520-T2** has the following indicator lights:



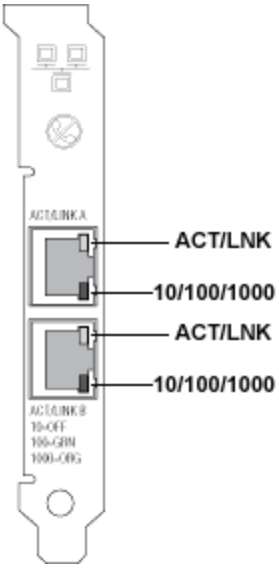
Label	Indication	Meaning
Link	Green	Linked at 10 Gb.
	Yellow	Linked at 1 Gb.
	Off	No link.
Activity	Blinking On/Off	Actively transmitting or receiving data.
	Off	No link.

The **Intel® Ethernet Server Adapter I350-T2, I340-T2, PRO/1000 P, PT Dual Port, and Gigabit ET Dual Port Server Adapters** have the following indicator lights:



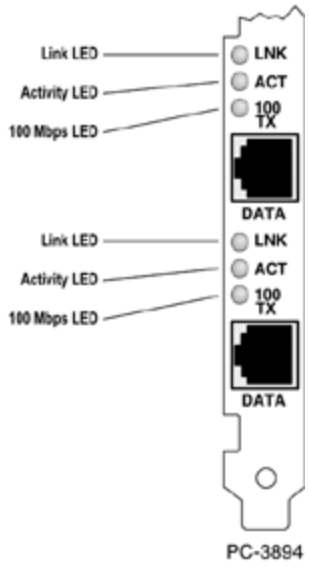
Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.
10/100/ 1000	Off	10 Mbps
	Green	100 Mbps
	Yellow	1000 Mbps
	Orange flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.

The **Intel® PRO/1000 MT and GT Dual Port Server Adapters** have the following indicator lights for each port:



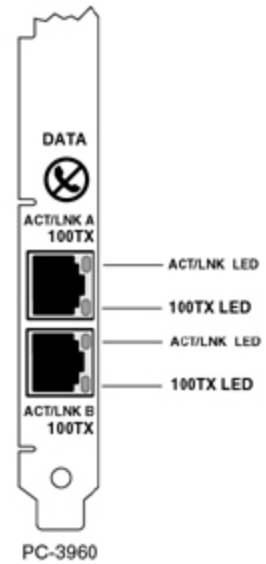
Label	Indication	Meaning
ACT/LNK	Green on	The port is connected to a valid link partner.
	Green flashing	Data activity
	Yellow flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
	Off	No link.
10=OFF 100=GRN 1000=ORG	Orange	1000 Mbps
	Green	100 Mbps
	Off	10 Mbps

The **PRO/100+ Dual Port Server adapter** (with three LEDs per port) has the following indicator lights:



Label	Indication	Meaning
LNK	On	The adapter and switch are receiving power; the cable connection between the switch and adapter is good.
	Off	The adapter and switch are not receiving power; the cable connection between the switch and adapter is faulty; or you have a driver configuration problem.
ACT	On or flashing	The adapter is sending or receiving network data. The frequency of the flashes varies with the amount of network traffic.
	Off	The adapter is not sending or receiving network data.
100	On	Operating at 100 Mbps.
	Off	Operating at 10 Mbps.

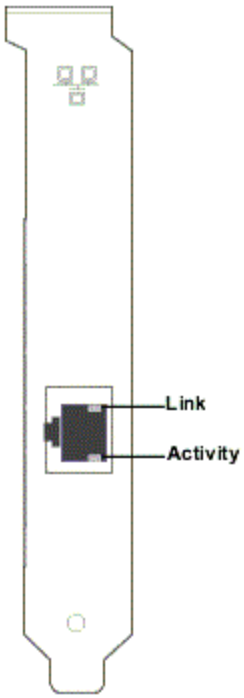
The **Intel® PRO/100 S Dual Port Server adapter** (with 64 bit PCI Connector) has the following indicator lights:



Label	Indication	Meaning
ACT/LNK	On	The adapter and switch are receiving power; the cable connection between the switch and adapter is good. NOTE: if only one port appears to be operating, go into the Network control panel (Windows) or a configuration file for other operating system, and check the status of the second port.
	Off	The adapter and switch are not receiving power, or you have a driver configuration problem.
	Flashing	The adapter is sending or receiving network data. The frequency of the flashes varies with the amount of network traffic.
100TX	On	Operating at 100 Mbps.
	Off	Operating at 10 Mbps.

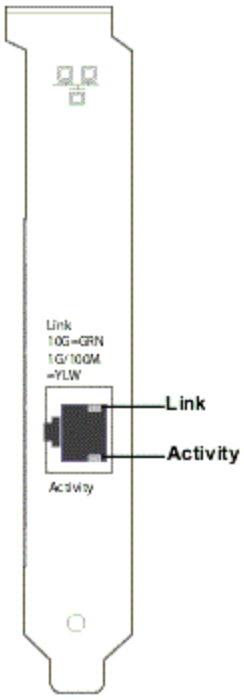
Single Port Copper Adapters

The Intel®**Ethernet Converged Network Adapter X550-T1** has the following indicator lights:



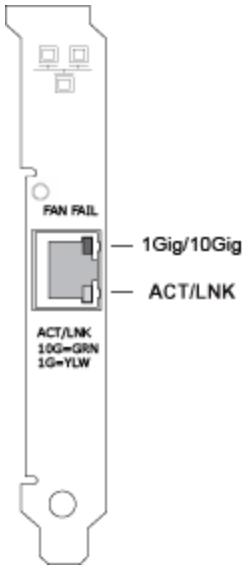
Label	Indication	Meaning
Link	Green	Linked at 10 Gb.
	Yellow	Linked at 1 Gb.
	Off	Linked at 100 Mbps.
Activity	Off	No link.
	Blinking On/Off	Actively transmitting or receiving data.

The Intel®**Ethernet Converged Network Adapter X540-T1** has the following indicator lights:



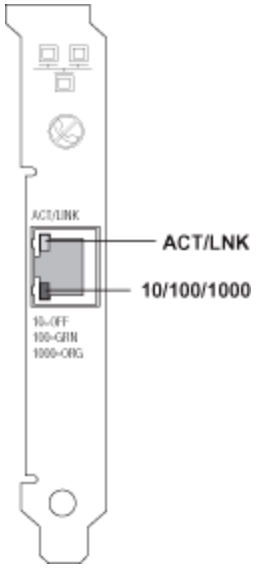
Label	Indication	Meaning
Link	Off	No link.
	Green	Linked at 10 Gb
	Yellow	Linked at 1 Gb
Activity	Off	No link.
	Blinking On/Off	Actively transmitting or receiving data.

The Intel® **10 Gigabit AT Server Adapter** has the following indicator lights:



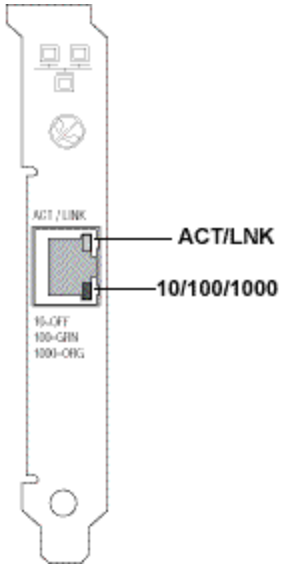
Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green blinking	Data activity
	Off	No link
	Green blinking	Identity. Use the "Identify Adapter" button in Intel PROSet or the BLINK command in IBAUtil to control blinking. See Intel PROSet Help for more information.
1Gig/10Gig	Off	No link
	Green	10 Gbps
	Yellow	1 Gbps
FAN FAIL	Red	Over temperature event occurred. The adapter will be shut down to avoid damage.

The Intel® **PRO/1000 PT Server Adapter** has the following indicator lights:



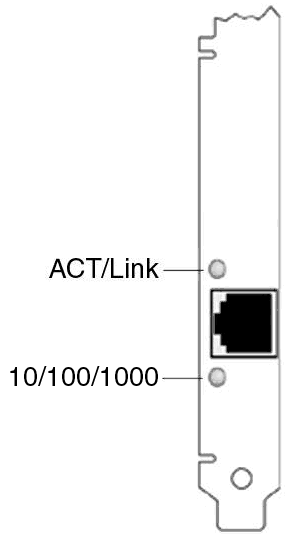
Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.
10=OFF 100=GRN 1000=ORG	Off	10 Mbps
	Green	100 Mbps
	Orange	1000 Mbps
	Orange flashing	Identity. Use the "Identify Adapter" button in Intel® PROSet to control blinking. See Intel PROSet Help for more information.

The Intel® **PRO/1000 MT Server Adapter** has the following indicator lights:



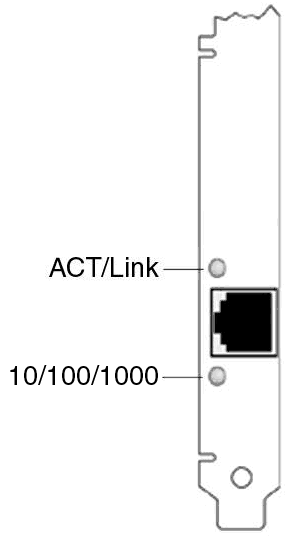
Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.
	Yellow flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
10=OFF 100=GRN 1000=ORG	Off	10 Mbps
	Green	100 Mbps
	Orange	1000 Mbps

The Intel® **Gigabit CT2, Gigabit CT, PRO/1000 T, and PRO/1000 MT Desktop Adapters** have the following indicator lights:



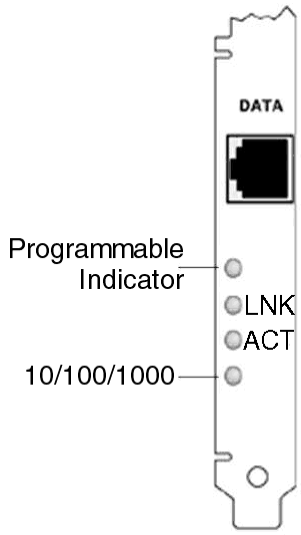
Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.
	Yellow flashing	Identity. Use the "Identify Adapter" button in Intel® PROSet to control blinking. See Intel PROSet Help for more information.
10/100/1000	Off	10 Mbps
	Green	100 Mbps
	Yellow	1000 Mbps

The Intel® **PRO/1000 XT Server Adapter** has the following indicator lights:



Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.
	Yellow flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
10/100/1000	Off	10 Mbps
	Green	100 Mbps
	Yellow	1000 Mbps

The Intel® **PRO/1000 T Server Adapter** has the following indicator lights:



Label	Indication	Meaning
not labeled	Flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
LNK	On	The adapter is connected to a valid link partner.
	Off	No link.
ACT	On	Data is being transmitted or received.
	Off	No data activity.
10/100/1000	Off	10 Mbps
	Green	100 Mbps
	Yellow	1000 Mbps

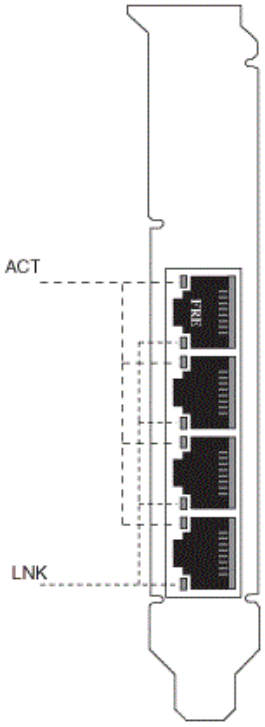
The Intel® **PRO/100+, PRO/100 M, PRO/100 S , and PRO/100 VE and VM Desktop adapters and Network Connections** have the following indicator lights:



Label	Indication	Meaning
ACT/LNK	On	The adapter and switch are receiving power; the cable connection between the switch and adapter is good.
	Off	The adapter and switch are not receiving power, or you have a driver configuration problem.
	Flashing	The adapter is sending or receiving network data. The frequency of the flashes varies with the amount of network traffic.
100TX	On	Operating at 100 Mbps.
	Off	Operating at 10 Mbps.

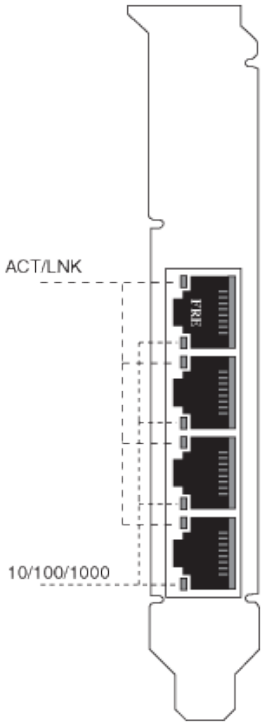
Quad Port Copper Adapters

The **Intel® Ethernet Converged Network Adapter X710-T4** has the following indicator lights:



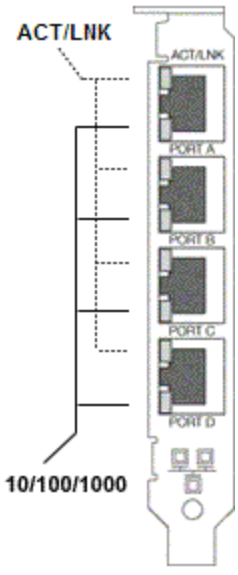
Label	Indication	Meaning
ACT	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.
LNK	Green	10 Gbps
	Yellow	1 Gbps
	Off	100 Mbps

The **Intel® Ethernet Server Adapter I350-T4, I340-T4, Gigabit ET and PRO/1000 PT Quad Port LP Server Adapters** have the following indicator lights:



Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.
10/100/ 1000	Green	100 Mbps
	Yellow	1000 Mbps
	Orange flashing	Identity. Use the "Identify Adapter" button in Intel® PROSet to control blinking. See Intel PROSet Help for more information.
	Off	10 Mbps

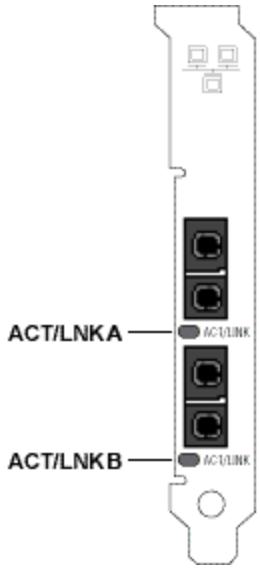
The Intel® **PRO/1000 MT, GT and PT Quad Port Server Adapters** have the following indicator lights for each port:



Label	Indication	Meaning
Top LED ACT/LNK	Green on	The port is connected to a valid link partner.
	Green flashing	Data activity
	Orange flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
	Off	No link.
Bottom LED 10/100/ 1000	Orange	1000 Mbps
	Green	100 Mbps
	Off	10 Mbps

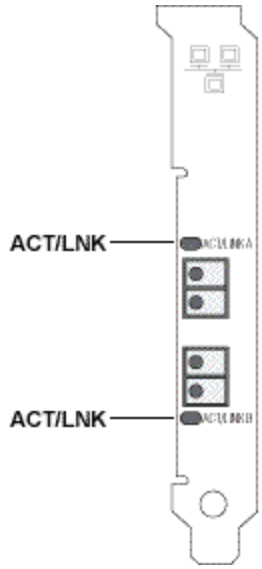
Dual Port Fiber Adapters

The Intel® **10 Gigabit XF SR Dual Port Server Adapters** has the following indicator lights:



Label	Indication	Meaning
ACT/LNK	On	The adapter is connected to a valid link partner.
	Blinking	Adapter is actively passing traffic. Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
	Off	No link.

The Intel® **PRO/1000 MF, PF, and Gigabit EF Dual Port Server Adapters** have the following indicator lights for each port:

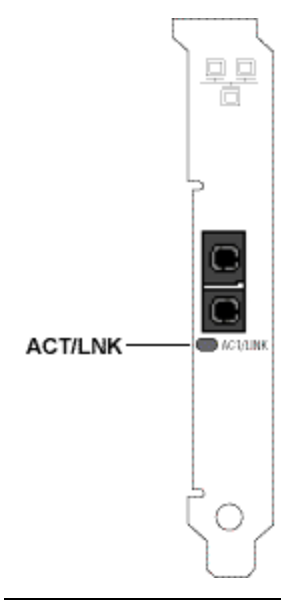


Label	Indication	Meaning
ACT/LNK	On	The adapter is connected to a valid link partner.
	Blinking	Adapter is actively passing traffic. Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
	Off	No link.

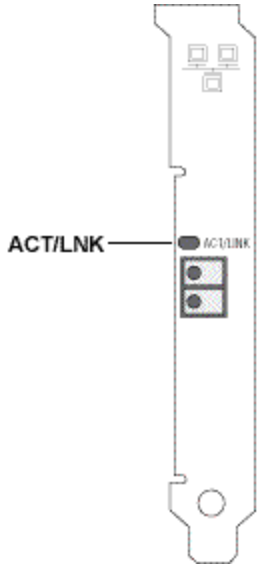
Single Port Fiber Adapters

The Intel® **10 Gigabit XF SR and LR Server Adapters** has the following indicator lights:

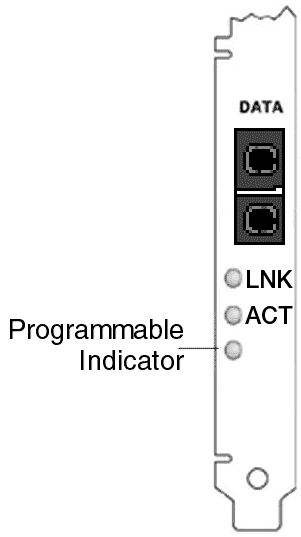
	Label	Indication	Meaning

			
	ACT/LNK	On	The adapter is connected to a valid link partner.
		Blinking	Adapter is actively passing traffic. Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
		Off	No link.

The Intel® **PRO/1000 MF and PF Server Adapters** have the following indicator lights:

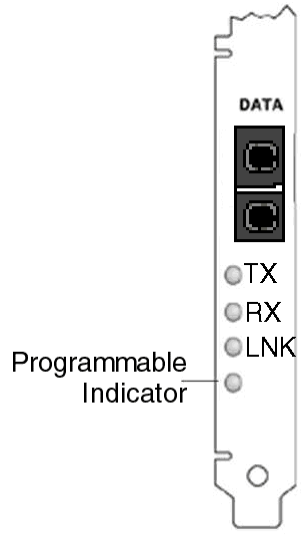
	Label	**Indication**	**Meaning**
			
	ACT/LNK	On	The adapter is connected to a valid link partner.
		Blinking	Adapter is actively passing traffic. Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.
		Off	No link.

The Intel® **PRO/1000 XF Server Adapter** has the following indicator lights:



Label	Indication	Meaning
LNK	On	The adapter is connected to a valid link partner.
	Off	No link.
ACT	On	Data is being transmitted or received.
	Off	No data activity.
not labeled	Flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.

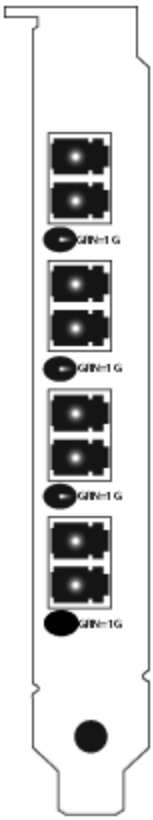
The Intel® **PRO/1000F Server adapter** has the following indicator lights:



Label	Indication	Meaning
TX	On	The adapter is sending data.
RX	On	The adapter is receiving data.
LNK	On	The adapter is connected to a valid link partner and is receiving link pulses.
	Off	No link.
not labeled	Flashing	Identity. Use the "Identify Adapter" button in Intel PROSet to control blinking. See Intel PROSet Help for more information.

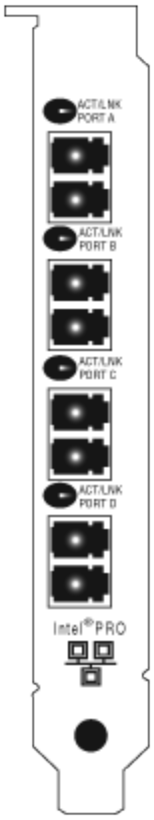
Quad Port Fiber Adapters

The Intel® **Ethernet Server Adapter I340-F4** has the following indicator lights:



Label	Indication	Meaning
GRN=1G	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.

The Intel® **PRO/1000 PF Quad Port Server Adapter** has the following indicator lights:



Label	Indication	Meaning
ACT/LNK	Green on	The adapter is connected to a valid link partner.
	Green flashing	Data activity
	Off	No link.

Known Issues

NOTE: iSCSI Known Issues and FCoE Known Issues are located in their own sections of this manual.

Lost Data Packets caused by Frequent LLDP Packets on an Inactive Port

When ports are teamed or bonded together in an active/passive configuration (for example, in a switch fault tolerance team, or a mode 1 bond), the inactive port may send out frequent LLDP packets, which results in lost data packets. This may occur with Intel ANS teaming on Microsoft Windows operating systems or with channel bonding on Linux systems. To resolve the issue, set one of the ports to be the Primary port.

Code 10 yellow bang errors on a Virtual Machine in Windows Device Manager

On a system running Microsoft Windows Server 2016, inside a Virtual Machine running Microsoft Windows Server 2016 or Windows Server 2012 R2, Intel Ethernet connections may have a code 10 yellow bang in Windows Device Manager. Installing a cumulative updated that contains Microsoft KB3192366 and KB3176936 will resolve the issue.

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 servers with several processors. Values greater than 4 may increase CPU utilization to unacceptable levels and have other negative impacts on system performance.

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.

Windows Known Issues

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.

Incomplete branding string displayed in the event log

Some branding strings are too long to be displayed fully in the event log. In these cases, the branding string will be truncated and the port's PCI Bus/Device/Function are appended to the string. For example: Intel(R) Ethernet Converged Network Ad... [129,0,1].

DCB QoS and Priority Flow Control do not act as expected

If you use Microsoft's Data Center Bridging (DCB) implementation configure Quality of Service (QoS) and Priority Flow Control (PFC), the actual traffic flow segregation per traffic class may not match your configuration and PFC may not pause traffic as expected. If you mapped more than one priority to a Traffic Class, enabling only one of the priorities and disabling the others will work around the issue. Installing Intel's DCB implementation will also resolve this issue. This issue affects Microsoft Windows Server 2012 and Server 2012 R2.

Link loss after changing the Jumbo Frames setting

Inside a guest partition on a Microsoft Windows Server 2012 R2 Hyper-V virtual machine, if you change the jumbo frame Advanced setting on an Intel® X540 based Ethernet Device or associated Hyper-V NetAdapter, you may lose link. Changing any other Advanced Setting will resolve the issue.

Virtual Machine Queues are not allocated until reboot

On a Microsoft Windows Server 2012 R2 system with Intel® Ethernet Gigabit Server adapters installed, if you install Hyper-V and create a VM switch, Virtual Machine Queues (VMQ) are not allocated until you reboot the system. Virtual machines can send and receive traffic on the default queue, but no VMQs will be used until after a system reboot.

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.

"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.

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.

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.

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.

Receive Side Scaling value is blank

Changing the Receive Side Scaling setting of an adapter in a team may cause the value for that setting to appear blank when you next check it. It may also appear blank for the other adapters in the team. The adapter may be unbound from the team in this situation. Disabling and enabling the team will resolve the issue.

RSS Load Balancing Profile Advanced Setting

Setting the "RSS load balancing profile" Advanced Setting to "ClosestProcessor" may significantly reduce CPU utilization. However, in some system configurations (such as a system with more Ethernet ports than processor cores), the "ClosestProcessor" setting may cause transmit and receive failures. Changing the setting to "NUMAScalingStatic" will resolve the issue.

Opening Windows Device Manager property sheet takes longer than expected

The Windows Device Manager property sheet may take 60 seconds or longer to open. The driver must discover all Intel Ethernet devices and initialize them before it can open the property sheet. This data is cached, so subsequent openings of the property sheet are generally quicker.

Power Management Known Issues

System does not wake when expected

Under Microsoft Windows Server 2012, the system may not wake even though Wake on LAN settings are enabled. Disabling Fast Startup in the operating system should resolve the issue.

System does not wake on link

On a driver-only installation, if you change 'Wake on Link Settings' to Forced and change 'Wake on Magic Packet' and 'Wake on Pattern Match' to Disabled, the system may not wake up when expected. In order to "Wake on Link" successfully, check the Power Management tab and make sure that “Allow this device to wake the computer” is checked. You may also need to change 'Wake on Magic Packet' or 'Wake on Pattern Match' to Enabled.

Directed Packets may not wake the system

On some systems, quad port server adapters may not wake when configured to wake on directed packet. If you experience problems waking on directed packets, you must configure the adapter to use Magic Packets*.

Power Management options are unavailable or missing

If you install only the base drivers, later install Intel® PROSet for Windows Device Manager, then remove Intel PROSet, the settings on the Power Management tab on the Adapter Property Sheet may be unavailable or missing altogether. You must reinstall Intel PROSet to resolve the issue.

System Wakes-Up from a Removed VLAN

If a system goes into standby mode, and a directed packet is sent to the IP address of the removed VLAN, the system will wake-up. This occurs because a directed packet bypasses VLAN filtering.

Intel Adapters ignore consecutive Wake Up signals while transitioning into standby mode

While sending a system into standby, occasionally a wake up packet arrives before the system completes the transition into standby mode. When this happens, the system ignores consecutive wake up signals and remains in standby mode until manually powered up using the mouse, keyboard, or power button.

Other Intel 10GbE Network Adapter Known Issues

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.

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® 10GbE Device. Doing so will cause loss of link.

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® 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.

Continuous PFC pause frames sent from Intel® Ethernet X520 based devices

If you have an Intel® Ethernet X520 based device connected to a switch port and modify the DCB bandwidth settings on the switch port, the Intel® Ethernet X520 device may perpetually send pause frames, causing a storm, and fail to transfer data to and from the storage targets it was using. To recover from this issue, disable the X520 ports, re-enable them, and then reconnect to the iSCSI target volumes. To avoid the issue, if the DCB bandwidth settings need to be changed, do one of the following:

- Power down the server that contains the Intel® Ethernet X520 device prior to modifying the DCB bandwidth settings.
- Disable the switch ports connected to Intel X520 based device.
- Have no traffic running on the Intel X520 based device.

Unexpected NMI with 82599-based NICs

If you set the PCIe Maximum Payload Size to 256 bytes in your system BIOS and install an 82599-based NIC, you may receive an NMI when the NIC attains link. This happens when the physical slot does not support a payload size of 256 Bytes even if the BIOS does. Moving the adapter to a slot that supports 256 bytes will resolve the issue. Consult your system documentation for information on supported payload values.

Intel® 710 Series Network Controller Known Issues

Some Intel® X710 based devices report a subvendor ID of 0x0000 and may display a generic branding string. Port 0 reports the correct subvendor ID and displays the correct branding string.

Intel X710 based devices may maintain link on any and all ports as long as power is provided to the device, regardless of the device's or system’s power state.

Intel X710/XL710 based devices do not support Microsoft Windows Server 2008 R2 or Windows Server 2012 as SR-IOV Guest operating systems. These devices do not support Windows Server 2012 as the SR-IOV Host operating system.

Driver Buffer Overflow Fix

The fix to resolve CVE-2016-8105, referenced in Intel SA-00069 https://security-center.intel.com/advisory.aspx?intelid=INTEL-SA-00069&languageid=en-fr, is included in this and future versions of the driver.

Intel ANS VLANs adversely affect performance

Intel ANS VLANs adversely affect the performance of X710 based devices. Use the networking features built into Microsoft Windows Server 2012, or other server management software, to assign VLANs.

Diagnostic button disabled

When part of an Intel® ANS team, devices based on the Intel X710 controller do not support diagnostics. Support will be added in a future release.

Unexpected IntelDCB errors in the Windows Application Event Log

After upgrading your X710 drivers, you may see several IntelDCB errors in the Windows Application Event Log. These errors are erroneous and can be ignored.

Lower than expected throughput on X710/XL710 based devices

If you have an X710 or XL710 based device installed in a four CPU socket system. Receive and transmit traffic may be significantly lower than expected. Setting your interrupt rate to High may mitigate the issue.

Quad Port Server Adapter Known Issues

Downshifting

When connecting to any Gigabit switch via a faulty CAT 5 cable where one pair is broken, the adapter does not downshift from 1 Gig to 100Mbps. For the adapter to downshift, it must identify two broken pairs in the cable.

System does not boot

Your system may run out of I/O resources and fail to boot if you install more than four quad port server adapters. Moving the adapters to different slots or rebalancing resources in the system BIOS may resolve the issue. This issue affects the following Adapters:

- Intel® Gigabit 4P I350-t Adapter

Regulatory Compliance Statements

FCC Class A Products

- Intel® Ethernet Network Adapter XXV710
- Intel® Ethernet Network Adapter XXV710-1
- Intel® Ethernet Network Adapter XXV710-2
- Intel® Ethernet I/O Module XL710-Q1
- Intel® Ethernet I/O Module XL710-Q2
- Intel® Ethernet Server Adapter X550-T2 for OCP
- Intel® Ethernet Server Adapter X550-T1 for OCP
- Intel® Ethernet Server Bypass Adapter X540-T2
- Intel® Ethernet Converged Network Adapter X540-T2
- Intel® Ethernet Converged Network Adapter X540-T1
- Intel® Ethernet OCP Server Adapter X520-1
- Intel® Ethernet OCP Server Adapter X520-2
- Intel® Ethernet Server Bypass Adapter X520-SR2
- Intel® Ethernet Server Adapter X520-T2
- Intel® Ethernet X520 10GbE Dual Port KX4 Mezz
- Intel® 10 Gigabit AT2 Server Adapter
- Intel® 10 Gigabit SR Dual Port Express Module
- Intel® 10 Gigabit AT Server Adapter
- Intel® Gigabit ET Quad Port Mezzanine Card
- Intel® PRO/1000 MT Quad Port Server Adapter

FCC Class B Products

- Intel® Ethernet Converged Network Adapter X710-2
- Intel® Ethernet Converged Network Adapter X710-4
- Intel® Ethernet Converged Network Adapter X710-T4
- Intel® Ethernet Converged Network Adapter XL710-Q1
- Intel® Ethernet Converged Network Adapter XL710-Q2
- Intel® Ethernet Converged Network Adapter X550-T1
- Intel® Ethernet Converged Network Adapter X550-T2
- Intel® Ethernet Server Adapter X520-1
- Intel® Ethernet Server Adapter X520-2
- Intel® Ethernet SFP+ LR Optics RMN FTLX1471
- Intel® Ethernet Converged Network Adapter X520-4
- Intel® Ethernet Converged Network Adapter X520-Q1
- Intel® 10 Gigabit AF DA Dual Port Server Adapter
- Intel® 10 Gigabit CX4 Dual Port Server Adapter
- Intel® 10 Gigabit XF SR Dual Port Server Adapter
- Intel® 10 Gigabit XF SR Server Adapter
- Intel® 10 Gigabit XF LR Server Adapter
- Intel® Ethernet Server Adapter I340-T2
- Intel® Ethernet Server Adapter I340-T4
- Intel® Ethernet Server Adapter I350-T2
- Intel® Ethernet Server Adapter I350-T4
- Intel® Ethernet Server Adapter I210-T1
- Intel® Ethernet Server Adapter I340-F4
- Intel® Gigabit ET Dual Port Server Adapter
- Intel® Gigabit ET Quad Port Server Adapter
- Intel® Gigabit EF Dual Port Server Adapter
- Intel® Gigabit CT Desktop Adapter
- Intel® PRO/1000 MT Server Adapter
- Intel® PRO/1000 MT Dual Port Server Adapter
- Intel® PRO/1000 GT Dual Port Server Adapter
- Intel® PRO/1000 PT Server Adapter
- Intel® PRO/1000 PT Dual Port Server Adapter
- Intel® PRO/1000 GT Quad Port Server Adapter
- Intel® PRO/1000 PT Quad Port Server Adapter
- Intel® PRO/1000 PT Quad Port LP Server Adapter
- Intel® PRO/1000 GT Desktop Adapter
- Intel® PRO/1000 PT Desktop Adapter
- Intel® PRO/1000 MF Server Adapter

- Intel® PRO/1000 MF Dual Port Server Adapter
- Intel® PRO/1000 PF Server Adapter
- Intel® PRO/1000 PF Dual Port Server Adapter
- Intel® PRO/1000 PF Quad Port Server Adapter
- Intel® PRO/100 M Desktop Adapter
- Intel® PRO/100 S Desktop Adapter
- Intel® PRO/100 S Server Adapter
- Intel® PRO/100 S Dual Port Server Adapter

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。

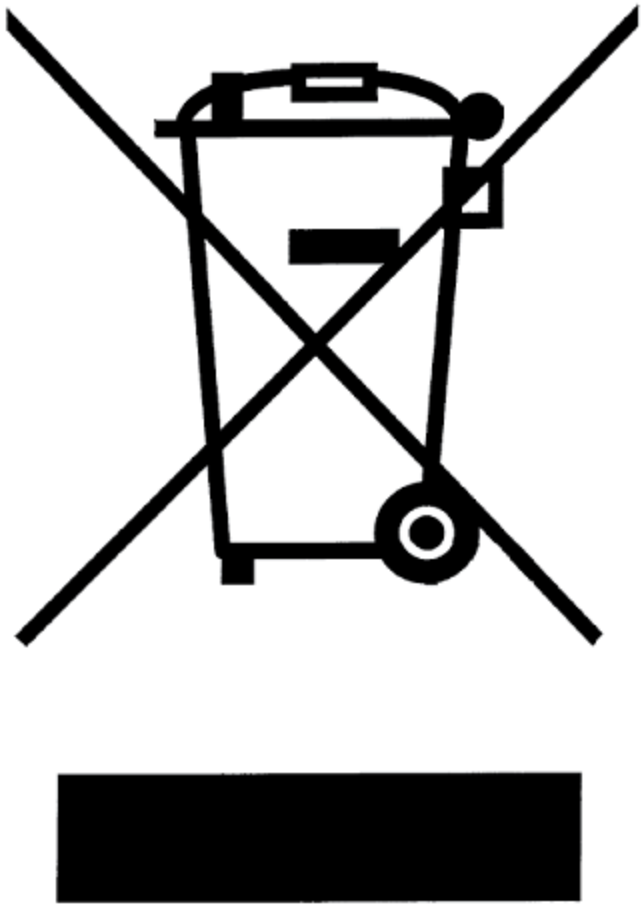
取扱説明書に従って正しい取り扱いをして下さい。

ＶＣＣＩ－Ｂ

KCC Notice Class B (Republic of Korea Only)

B급 기기 (가정용 방송통신기기)	이 기기는 가정용(B급)으로 전자파적합등록을 한 기기로서 주로 가정에서 사용하는 것을 목적으로 하며, 모든 지역에서 사용할 수 있습니다.
CLASS B device residential broadcasting and communication equipment	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.

Customer Support

Intel support is available on the web or by phone. Support offers the most up-to-date information about Intel products, including installation instructions, troubleshooting tips, and general product information.

Web and Internet Sites

Support: http://www.intel.com/support

Corporate Site for Network Products: http://www.intel.com/products/ethernet/overview.htm

Legal Disclaimers

INTEL SOFTWARE LICENSE AGREEMENT

IMPORTANT - READ BEFORE COPYING, INSTALLING OR USING.

Do not copy, install, or use this software and any associated materials (collectively, the "Software") provided under this license agreement ("Agreement") until you have carefully read the following terms and conditions.

By copying, installing, or otherwise using the Software, you agree to be bound by the terms of this Agreement. If you do not agree to the terms of this Agreement, do not copy, install, or use the Software.

LICENSES

Please Note:

- If you are a network or system administrator, the "Site License" below shall apply to you.
- If you are an end user, the "Single User License" shall apply to you.
- If you are an original equipment manufacturer (OEM), the "OEM 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 (a) physical Intel component products, and (b) virtual ("emulated") devices designed to appear as Intel component products to a Guest operating system running within the context of a virtual machine. Any other use of the Software, including but not limited to use with non-Intel component products, is not licensed hereunder.**
2. Subject to all of the terms and conditions of this Agreement, Intel Corporation ("Intel") grants to you a non-exclusive, non-assignable, copyright license to use the Software.
3. 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.
4. You may not reverse engineer, decompile, or disassemble the Software.
5. The Software may include portions offered on terms differing from 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 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 (a) physical Intel component products, and (b) virtual ("emulated") devices designed to appear as Intel component products to a Guest operating system running within the context of a virtual machine. Any other use of the Software, including but not limited to use with non-Intel component products, is not licensed hereunder.**
2. Subject to all of the terms and conditions of this Agreement, Intel Corporation ("Intel") grants to you a non-exclusive, non-assignable, copyright license to use the Software.
3. 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.
4. You may not reverse engineer, decompile, or disassemble the Software.
5. The Software may include portions offered on terms differing from those set out here, as set out in a license accompanying those portions.

OEM LICENSE: You may reproduce and distribute the Software only as an integral part of or incorporated in your product, as a standalone Software maintenance update for existing end users of your products, excluding any other standalone products, or as a component of a larger Software distribution, including but not limited to the distribution of an installation image or a Guest Virtual Machine image, subject to these conditions:

1. **This Software is licensed for use only in conjunction with (a) physical Intel component products, and (b) virtual ("emulated") devices designed to appear as Intel component products to a Guest operating system running within the context of a virtual machine. Any other use of the Software, including but not limited to use with non-Intel component products, is not licensed hereunder.**
2. Subject to all of the terms and conditions of this Agreement, Intel Corporation ("Intel") grants to you a non-exclusive, non-assignable, copyright license to use the Software.
3. 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.
4. You may not reverse engineer, decompile, or disassemble the Software.
5. You may only distribute the Software to your customers pursuant to a written license agreement. Such license agreement may be a "break-the-seal" license agreement. At a minimum such license shall safe-guard Intel's ownership rights to the Software.
6. You may not distribute, sublicense or transfer the Source Code form of any components of the Software and derivatives thereof to any third party without the express written consent of Intel.
7. The Software may include portions offered on terms differing from those set out here, as set out in a license accompanying those portions.

LICENSE RESTRICTIONS. You may NOT: (i) use or copy the Software except as provided in this Agreement; (ii) rent or lease the Software to any third party; (iii) assign this Agreement or transfer the Software without the express written consent of Intel; (iv) modify, adapt, or translate the Software in whole or in part except as provided in this Agreement; (v) reverse engineer, decompile, or disassemble the Software; (vi) attempt to modify or tamper with the normal function of a license manager that regulates usage of the Software; (vii) distribute, sublicense or transfer the Source Code form of any components of the Software and derivatives thereof to any third party without the express written consent of Intel; (viii) permit, authorize, license or sublicense any third party to view or use the Source Code; (ix) modify or distribute the Source Code or Software so that any part of it becomes subject to an Excluded License. (An "Excluded License" is one that requires, as a condition of use, modification, or distribution, that (a) the code be disclosed or distributed in source code form; or (b) others have the right to modify it.); (x) use or include the Source Code or Software in deceptive, malicious or unlawful programs.

NO OTHER RIGHTS. No rights or licenses are granted by Intel to you, expressly or by implication, with respect to any proprietary information or patent, copyright, mask work, trademark, trade secret, or other intellectual property right owned or controlled by Intel, except as expressly provided in this Agreement. Except as expressly provided herein, no license or right is granted to you directly or by implication, inducement, estoppel, or otherwise. Specifically, Intel grants no express or implied right to you under Intel patents, copyrights, trademarks, or other intellectual property rights.

OWNERSHIP OF SOFTWARE AND COPYRIGHTS. The Software is licensed, not sold. Title to all copies of the Software remains with Intel. 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. You agree to prevent any unauthorized copying of 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.

ADDITIONAL TERMS FOR PRE-RELEASE SOFTWARE. If the Software you are installing or using under this Agreement is pre-commercial release or is labeled or otherwise represented as "alpha-" or "beta-" versions of the Software ("pre-release Software"), then the following terms apply. To the extent that any provision in this Section conflicts with any other term(s) or condition(s) in this Agreement with respect to pre-release Software, this Section shall supersede the other term(s) or condition(s), but only to the extent necessary to resolve the conflict. You understand and acknowledge that the Software is pre-release Software, does not represent the final

Software from Intel, and may contain errors and other problems that could cause data loss, system failures, or other errors. The pre-release Software is provided to you "as-is" and Intel disclaims any warranty or liability to you for any damages that arise out of the use of the pre-release Software. You acknowledge that Intel has not promised that pre-release Software will be released in the future, that Intel has no express or implied obligation to you to release the pre-release Software and that Intel may not introduce Software that is compatible with the pre-release Software. You acknowledge that the entirety of any research or development you perform that is related to the pre-release Software or to any product making use of or associated with the pre-release Software is done at your own risk. If Intel has provided you with pre-release Software pursuant to a separate written agreement, your use of the pre-release Software is also governed by such agreement.

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. In the event that you use the Software in conjunction with a virtual ("emulated") device designed to appear as an Intel component product, you acknowledge that Intel is neither the author nor the creator of the virtual ("emulated") device. You understand and acknowledge that Intel makes no representations about the correct operation of the Software when used with a virtual ("emulated") device, that Intel did not design the Software to operate in conjunction with the virtual ("emulated") device, and that the Software may not be capable of correct operation in conjunction with the virtual ("emulated") device. You agree to assume the risk that the Software may not operate properly in conjunction with the virtual ("emulated") device. You agree to indemnify and hold Intel and its officers, subsidiaries and affiliates harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of product liability, personal injury or death associated with the use of the Software in conjunction with the virtual ("emulated") device, even if such claim alleges that Intel was negligent regarding the design or manufacture of the Software.

UNAUTHORIZED USE.THE SOFTWARE IS NOT DESIGNED, INTENDED, OR AUTHORIZED FOR USE IN ANY TYPE OF SYSTEM OR APPLICATION IN WHICH THE FAILURE OF THE SOFTWARE COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR (E.G MEDICAL SYSTEMS, LIFE SUSTAINING OR LIFE SAVING SYSTEMS). If you use the Software for any such unintended or unauthorized use, you shall indemnify and hold Intel and its officers, subsidiaries and affiliates harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of product liability, personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Intel was negligent regarding the design or manufacture of the part.

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 the State of California, without regard to principles of conflict of laws. You agree that the terms of the United Nations Convention on Contracts for the Sale of Goods do not apply to this Agreement. 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 enclosed Software and documentation were developed at private expense, and are provided with "RESTRICTED RIGHTS." Use, duplication, or disclosure by the Government is subject to restrictions as set forth in FAR 52.227-14 and DFARS 252.227-7013 et seq. or its successor. The use of this product by the Government constitutes acknowledgment of Intel's proprietary rights in the Software. Contractor or Manufacturer is Intel.

LANGUAGE; TRANSLATIONS. In the event that the English language version of this Agreement is accompanied by any other version translated into any other language, such translated version is provided for convenience purposes only and the English language version shall control.

Third-party Licenses

Portions of this release may include software distributed under the following licenses.

Open Toolkit Library (OpenTK)

The Open Toolkit library license

Copyright (c) 2006 - 2009 The Open Toolkit library.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Third parties

* The Open Toolkit library includes portions of the Mono class library, which are covered by the following license:

Copyright (c) 2004 Novell, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR

COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

* Half-to-Single and Single-to-Half conversions are covered by the following license:

Copyright (c) 2002, Industrial Light & Magic, a division of Lucas Digital Ltd. LLC. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Industrial Light & Magic nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

RSA Data Security-MD5 Message

RSA Data Security

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

Copyright and Legal Disclaimers

Copyright © 2002-2018 Intel Corporation. All rights reserved.

Intel Corporation, 5200 N.E. Elam Young Parkway, Hillsboro, OR 97124-6497 USA

Intel Corporation assumes no responsibility for errors or omissions in this document. Nor does Intel make any commitment to update the information contained herein.

Intel, the Intel logo, and Itanium are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

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

Limited Lifetime Hardware Warranty

Intel warrants to the original owner that the adapter product delivered in this package will be free from defects in material and workmanship. This warranty does not cover the adapter product if it is damaged in the process of being installed or improperly used.

THE ABOVE WARRANTY IS IN LIEU OF ANY OTHER WARRANTY, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF NONINFRINGEMENT OF INTELLECTUAL PROPERTY, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE.

This warranty does not cover replacement of adapter products damaged by abuse, accident, misuse, neglect, alteration, repair, disaster, improper installation, or improper testing. If the adapter product is found to be defective, Intel, at its option, will replace or repair the hardware product at no charge except as set forth below, or refund your purchase price provided that you deliver the adapter product along with a Return Material Authorization(RMA) number (see below), along with proof of purchase (if not registered), either to the dealer from whom you purchased it or to Intel with an explanation of any deficiency. If you ship the adapter product, you must assume the risk of damage or loss in transit. You must use the original container (or the equivalent) and pay the shipping charge.

Intel may replace or repair the adapter product with either new or reconditioned parts, and any adapter product, or part thereof replaced by Intel becomes Intel's property. Repaired or replaced adapter products will be returned to you at the same revision level as received or higher, at Intel's option. Intel reserves the right to replace discontinued adapter products with an equivalent current generation adapter product.

Returning a defective product

From North America:

Before returning any adapter product, contact Intel Customer Support and obtain a Return Material Authorization (RMA) number by calling +1 916-377-7000.

If the Customer Support Group verifies that the adapter product is defective, they will have the RMA department issue you an RMA number to place on the outer package of the adapter product. Intel cannot accept any product without an RMA number on the package.

All other locations:

Return the adapter product to the place of purchase for a refund or replacement.

Intel Adapter Money-back Guarantee (North America Only)

Intel wants you to be completely satisfied with the Intel adapter product that you have purchased. Any time within ninety (90) days of purchase, you may return your Intel adapter to the original place of purchase for a full refund of the purchase price from your dealer. Resellers and distributors, respectively, accepting returns and refunding money back to their customers may return Intel adapters to their original place of purchase. Intel guarantees that it will accept returns under this policy and refund the original purchase price to customers purchasing directly from Intel.

Limitation of Liability and Remedies

INTEL'S SOLE LIABILITY HEREUNDER SHALL BE LIMITED TO DIRECT, OBJECTIVELY MEASURABLE DAMAGES. IN NO EVENT SHALL INTEL HAVE ANY LIABILITY FOR ANY INDIRECT OR SPECULATIVE DAMAGES (INCLUDING, WITHOUT LIMITING THE FOREGOING, CONSEQUENTIAL, INCIDENTAL, AND SPECIAL DAMAGES) INCLUDING, BUT NOT LIMITED TO, INFRINGEMENT OF INTELLECTUAL PROPERTY, REPROCUREMENT COSTS, LOSS OF USE, BUSINESS INTERRUPTIONS, LOSS OF GOODWILL, AND LOSS OF PROFITS, WHETHER ANY SUCH DAMAGES ARISE OUT OF CONTRACT NEGLIGENCE, TORT, OR UNDER ANY WARRANTY, IRRESPECTIVE OF WHETHER INTEL HAS ADVANCE NOTICE OF THE POSSIBILITY OF ANY SUCH DAMAGES. NOTWITHSTANDING THE

FOREGOING, INTEL'S TOTAL LIABILITY FOR ALL CLAIMS UNDER THIS AGREEMENT SHALL NOT EXCEED THE PRICE PAID FOR THE PRODUCT. THESE LIMITATIONS ON POTENTIAL LIABILITIES WERE AN ESSENTIAL ELEMENT IN SETTING THE PRODUCT PRICE. INTEL NEITHER ASSUMES NOR AUTHORIZES ANYONE TO ASSUME FOR IT ANY OTHER LIABILITIES.

Some states do not allow the exclusion or limitation of incidental or consequential damages, so the above limitations may not apply to you.

Critical Control Applications: Intel specifically disclaims liability for use of the adapter product in critical control applications (including, for example only, safety or health care control systems, nuclear energy control systems, or air or ground traffic control systems) by Licensee or Sublicensees, and such use is entirely at the user's risk. Licensee agrees to defend, indemnify, and hold Intel harmless from and against any and all claims arising out of use of the adapter product in such applications by Licensee or Sublicensees.

Software: Software provided with the adapter product is not covered under the hardware warranty described above. See the applicable software license agreement which shipped with the adapter product for details on any software warranty.

Note: If you are a consumer under the Australian Consumer Law, this warranty does not apply to you. Please visit Australian Limited Lifetime Hardware Warranty to view the limited warranty which is applicable to Australian consumers.