Intel® Ethernet Adapters and Devices User Guide

Rev. 29.4

Contents

Welcome to Intel® Ethernet Adapters and Devices User Guide!

Welcome to the software user guide for Intel® Ethernet adapters and devices. This guide covers software installation, setup procedures, troubleshooting tips, and other information for Intel network adapters, Ethernet connections, and other devices based on the following:

- Intel® Ethernet 800 Series
- Intel® Ethernet 700 Series
- Intel® Ethernet 500 Series
- Intel® Ethernet 300 Series
- Intel® Ethernet 200 Series

What's New in This Release

- View the Intel® Ethernet Controller Products Release Notes on intel.com for the high-level list of what has changed in this release.
- See the Release Notes page in this user guide for information how to view the release notes for past releases.

Note:

Point releases (for example, 29.1.1) may not have a corresponding user guide version. In those cases, refer to the user guide for the major release (for example, 29.1).

Supported Operating Systems and Devices

Refer to the following pages for detailed information:

- Supported Operating Systems
- Supported Hardware

Intended Audience

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

Overview

This section covers basic information about the operating systems and hardware supported in this release. It also provides links to additional documentation.

- Supported Operating Systems

- Supported Hardware
- Related Documentation
- Discontinued Support

Supported Operating Systems

The drivers in this software release have been tested with the operating systems (OSs) listed below. Additional OSs may function with our drivers but are not tested.

> **Note:**
>
> Not all devices support all operating systems listed. Refer to the Release Notes for detailed OS support information for your device.
>
> Refer to Discontinued Support for information on the last release to support particular versions of an operating system.

Microsoft Windows Server*, Azure Stack HCI, and Windows*

- Microsoft Windows Server* 2025
- Microsoft Windows Server 2022
- Microsoft Windows Server 2019, Version 1903
- Microsoft Windows Server 2016
- Microsoft Azure Stack HCI
- Microsoft Windows* 11 version 24H2
- Microsoft Windows 11 version 23H2 (build 22631.2506)
- Microsoft Windows 11 version 22H2 (build 22621)
- Microsoft Windows 10 version 21H2 (build 19044)
- Microsoft Windows 10 RS5, Version 1809 (build 17763)

> **Note:**
>
> - Devices based on the following do not support Microsoft Windows or Windows Server:
>   - Intel® Ethernet Connection E822-C
>   - Intel® Ethernet Connection E822-L
> - Microsoft Windows 32-bit operating systems are only supported on Intel 1Gbps Ethernet Adapters.
> - 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.

- In Microsoft Windows Server 2025, all devices based on the Intel Ethernet 710 Series are reported with generic, 2-part device IDs and branding strings, and will be configured as generic devices with no custom default settings. This change resolves an issue seen when deploying network intent roles using the Network ATC scripts in Windows Server 2025 and Azure Stack HCI.

VMware ESXi*

- VMware ESXi* 8.0
- VMware ESXi 7.0

Please refer to VMware's download site for the latest ESXi drivers for Intel Ethernet devices.

Linux*

- Linux* Real Time Kernel 5.x and 4.x [1]
- Linux, v2.4 kernel or higher
- Red Hat Enterprise Linux* (RHEL) 9.4
- Red Hat Enterprise Linux 8.10
- SUSE Linux Enterprise Server* (SLES) 15 SP6
- SUSE Linux Enterprise Server 12 SP5
- Canonical Ubuntu* 24.04 LTS
- Canonical Ubuntu 22.04 LTS
- Debian* 11
- openEuler* 22.03 LTS SP3 for AArch64 [2]

[1] Only supported on Intel® Ethernet E810 Series.

[2] Only supported on Intel Ethernet E810 Series, with the Linux ice and iavf drivers and select Intel® Network Connection Tools. See tool readmes for details.

FreeBSD*

- FreeBSD* 14.1
- FreeBSD 13.3

Supported Hardware

This software release supports Intel® Ethernet devices based on the silicon controllers and connections listed below.

For help identifying your network device and finding supported devices, visit https://www.intel.com/support.

Note:

Available features and settings are dependent on your device and operating system. **Not all settings are available on every device/OS combination.**

Refer to the Release Notes and Feature Support Matrix documents for details on supported features and OS combinations.

Intel® Ethernet 800 Series

- Intel® Ethernet Controller E810-C
- Intel® Ethernet Controller E810-XXV
- Intel® Ethernet Connection E822-C
- Intel® Ethernet Connection E822-L
- Intel® Ethernet Connection E823-C
- Intel® Ethernet Connection E823-L

Intel® Ethernet 700 Series

- Intel® Ethernet Controller I710
- Intel® Ethernet Controller X710
- Intel® Ethernet Controller XL710
- Intel® Ethernet Network Connection X722
- Intel® Ethernet Controller XXV710
- Intel® Ethernet Controller V710

Intel® Ethernet 500 Series

- Intel® Ethernet Controller 82599
- Intel® Ethernet Controller X520
- Intel® Ethernet Controller X550
- Intel® Ethernet Controller X552
- Intel® Ethernet Controller X553

Intel® Ethernet 300 Series and Other

- Intel® I210 Gigabit Ethernet Controller
- Intel® I350 Gigabit Ethernet Controller
- Intel® Ethernet Controller I225
- Intel® Ethernet Controller I226

- Intel® Ethernet Connection I217
- Intel® Ethernet Connection I218
- Intel® Ethernet Connection I219

Compatibility Notes

In addition, note the following limitations for Intel Ethernet devices and connections.

Intel Ethernet 800 Series:

- Devices based on the Intel Ethernet Controller E810-C have an expected total throughput for the entire device of 100Gbps in each direction if one 100G cable is connected or if two 100G cables are connected.
- Devices based on the Intel Ethernet 800 Series do not support RDMA when operating in multiport mode with more than 4 ports.

Intel Ethernet 700 Series:

- Devices based on the Intel Ethernet Controller XL710 (4x10Gbps, 1x40Gbps, and 2x40Gbps) have an expected total throughput for the entire device of 40Gbps in each direction.
- The first port of Intel Ethernet 700 Series adapters will display the correct branding string. All other ports on the same device will display a generic branding string.
- In Microsoft Windows Server 2025, all devices based on the Intel Ethernet 710 Series are reported with generic, 2-part device IDs and branding strings, and will be configured as generic devices with no custom default settings. This change resolves an issue seen when deploying network intent roles using the Network ATC scripts in Windows Server 2025 and Azure Stack HCI.
- For an Intel Ethernet 700 Series 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 limit the throughput of the adapter.
- Devices based on the Intel Ethernet Controller X722 do not support the following features:
  - Intel® PROSet for Windows* Device Manager
  - Intel® Advanced Network Services (Intel® ANS) teams or VLANs (LBFO is supported)

Intel Ethernet 500 Series:

- Devices based on the Intel Ethernet Connection X552 and Intel Ethernet Connection X553 do not support the following features:
  - Energy Efficient Ethernet (EEE)
  - Intel PROSet for Windows Device Manager
  - Intel ANS teams or VLANs (LBFO is supported)
  - Data Center Bridging (DCB)
  - IPSec Offloading
  - MACSec Offloading

- In addition, SFP+ devices based on the Intel Ethernet Connection X552 and Intel Ethernet Connection X553 do not support the following features:
    - Speed and duplex auto-negotiation
    - Wake on LAN
    - 1000BASE-T SFP Modules

Related Documentation

Intel Resource and Documentation Center

Visit the Intel Resource and Documentation Center for additional resources, configuration guides, and technical documentation on Intel Ethernet software and devices:

- https://www.intel.com/content/www/us/en/resources-documentation/developer.html

You can find content related to Intel Ethernet in the following collection:

- https://www.intel.com/content/www/us/en/resources-documentation/developer.html?collection=for-hardware/ethernet-products

Note: Some content may require a login.

Release Notes and Feature Support Matrixes

Please see the Intel® Ethernet Controller Products Release Notes for details about new or removed features, supported operating systems per device family, and known issues and limitations. See the Release Notes page in this user guide for information how to view the release notes for the current or past releases.

Refer to the Feature Support Matrix for your device family for information on supported features, cables, media types, operating systems, and more.

Application Notes and Configuration Guides

Numerous app notes, configuration guides, and other software documentation are available in the Intel Ethernet collection on the Intel Resource and Documentation Center. You can view or search available documentation by device family.

For newer devices, guides may be available on the following topics:

- Performance tuning guides for Linux* or Microsoft Windows*
- Application Device Queues (ADQ)
- Dynamic Device Personalization (DDP)
- Switchdev mode
- Advanced Precision Time Protocol (PTP) and SyncE
- Remote Direct Memory Access (RDMA)

User Guides for Specific Devices

Some adapters and devices have user guides with detailed configuration and setup information. You can access public versions of these documents in the Intel Resource and Documentation Center at the links below.

Affected Products	Type of Document + Link
Intel® Ethernet Network Adapter E810-XXV-4T	User guide (PDF)
Intel® Ethernet Network Adapter E810-C-Q2T	User guide (PDF)

Discontinued Support

This page lists the features, devices, or operating systems that were discontinued in a particular release. Subsequent releases will not support the discontinued feature.

Release 28.3

End of support for:

- Microsoft Windows Server 2012 and Windows Server 2012 R2
- Microsoft Azure Stack HCI, version 21H2 and 20H2
- Intel® I354 Gigabit Ethernet Controller and the e1s driver
- Autorun.exe and its associated files

Release 28.2.3 is the last release that includes drivers for these OSs and devices.

Release 28.1

End of support for RSS on Microsoft Windows operating systems for the following devices:

- Intel® Ethernet Connection I217-LM
- Intel® Ethernet Connection I218-LM
- Intel® Ethernet Connection (2) I218-LM
- Intel® Ethernet Connection (3) I218-LM
- Intel® Ethernet Connection I218-LM
- Intel® Ethernet Connection I218-V

- Intel® Ethernet Connection (2) I218-V
- Intel® Ethernet Connection (3) I218-V
- Intel® Ethernet Connection I218-V
- Intel® Ethernet Connection I217-V
- Intel® Ethernet Network Adapter I226-T1
- Intel® Ethernet Network Adapter I226-T1
- Intel® Ethernet Network Adapter I225-T1
- Intel® Ethernet Network Adapter I225-T1
- Intel® Ethernet Controller (2) I225-IT
- Intel® Killer™ E3100X 2.5 Gigabit Ethernet Controller
- Intel® Ethernet Controller (2) I225-LM
- Intel® Ethernet Controller (2) I225-LMvP
- Intel® Ethernet Controller (2) I225-V
- Intel® Ethernet Controller (3) I225-IT
- Intel® Killer™ E3100X 2.5 Gigabit Ethernet Controller
- Intel® Ethernet Controller (3) I225-LM
- Intel® Ethernet Controller (3) I225-LMvP
- Intel® Ethernet Controller (3) I225-V
- Intel® Ethernet Controller I226-IT
- Intel® Killer™ E3100X 2.5 Gigabit Ethernet Controller (3)
- Intel® Ethernet Controller I226-LM
- Intel® Ethernet Controller I226-LMvP
- Intel® Ethernet Controller I226-V
- Intel® Ethernet Controller I225-LM
- Intel® Ethernet Controller I225-V
- Intel® Ethernet Connection (2) I219-LM
- Intel® Ethernet Connection (2) I219-V
- Intel® Ethernet Connection (3) I219-LM
- Intel® Ethernet Connection (4) I219-LM
- Intel® Ethernet Connection (4) I219-V
- Intel® Ethernet Connection (5) I219-LM
- Intel® Ethernet Connection (5) I219-V
- Intel® Ethernet Connection I219-LM
- Intel® Ethernet Connection (10) I219-LM
- Intel® Ethernet Connection (11) I219-LM
- Intel® Ethernet Connection (12) I219-LM

- Intel® Ethernet Connection (13) I219-LM
- Intel® Ethernet Connection (14) I219-LM
- Intel® Ethernet Connection (15) I219-LM
- Intel® Ethernet Connection (16) I219-LM
- Intel® Ethernet Connection (17) I219-LM
- Intel® Ethernet Connection (18) I219-LM
- Intel® Ethernet Connection (19) I219-LM
- Intel® Ethernet Connection (20) I219-LM
- Intel® Ethernet Connection (22) I219-LM
- Intel® Ethernet Connection (23) I219-LM
- Intel® Ethernet Connection (24) I219-LM
- Intel® Ethernet Connection (25) I219-LM
- Intel® Ethernet Connection (26) I219-LM
- Intel® Ethernet Connection (27) I219-LM
- Intel® Ethernet Connection (6) I219-LM
- Intel® Ethernet Connection (7) I219-LM
- Intel® Ethernet Connection (8) I219-LM
- Intel® Ethernet Connection (9) I219-LM
- Intel® Ethernet Connection I219-V
- Intel® Ethernet Connection (10) I219-V
- Intel® Ethernet Connection (11) I219-V
- Intel® Ethernet Connection (12) I219-V
- Intel® Ethernet Connection (13) I219-V
- Intel® Ethernet Connection (14) I219-V
- Intel® Ethernet Connection (16) I219-V
- Intel® Ethernet Connection (17) I219-V
- Intel® Ethernet Connection (18) I219-V
- Intel® Ethernet Connection (19) I219-V
- Intel® Ethernet Connection (20) I219-V
- Intel® Ethernet Connection (22) I219-V
- Intel® Ethernet Connection (23) I219-V
- Intel® Ethernet Connection (24) I219-V
- Intel® Ethernet Connection (25) I219-V
- Intel® Ethernet Connection (26) I219-V
- Intel® Ethernet Connection (27) I219-V

- Intel® Ethernet Connection (6) I219-V
- Intel® Ethernet Connection (7) I219-V
- Intel® Ethernet Connection (8) I219-V
- Intel® Ethernet Connection (9) I219-V
- Intel® Ethernet Connection (3) I218-LM
- Intel® Ethernet Connection (3) I218-V

Release 27.9

Release 27.9 is the last release that supports the following:

- Microsoft Windows 8.1
- Microsoft Windows 10 Version 20H2 (build 19042)
- Microsoft Windows 10 Version 21H1 (build 19043)
- VMWare ESXi 6.5
- VMWare ESXi 6.7
- iw_ixl FreeBSD driver

Release 27.5

Release 27.5 is the last release that includes the Microsoft e1q driver in your download package. Release 27.6 removed the e1q driver from the installation package. This affects the following devices:

- Intel® 82575EB Gigabit Network Connection
- Intel® 82575EB Gigabit Backplane Connection
- Intel® Gigabit VT Quad Port Server Adapter
- Intel® 82575EB Multi-Function Network Device
- Intel® 82574L Gigabit Network Connection
- Intel® 82583V Gigabit Network Connection
- Intel® Gigabit CT Desktop Adapter
- Intel® Gigabit CT2 Desktop Adapter
- Intel® 82576 Gigabit Dual Port Network Connection
- Intel® Gigabit ET Dual Port Server Adapter
- Intel® 82576NS Gigabit Ethernet Controller
- Intel® 82576NS Gigabit Network Connection
- Intel® 82576NS SerDes Gigabit Ethernet Controller
- Intel® Gigabit EF Dual Port Server Adapter
- Intel® 82576 Gigabit Dual Port Server Network Connection
- Intel® Gigabit ET Quad Port Server Adapter

- Intel® Gigabit ET2 Quad Port Server Adapter
- Intel® 82576 Gigabit Dual Port Network Connection
- Intel® Gigabit ET Quad Port Mezzanine Card

Release 27.0

Release 27.0 the last release that supports:

- The Microsoft Windows v1q driver

This driver will no longer be tested or updated. The driver may still be provided in your download package or on your install media for your convenience.

Release 26.7

Release 26.7 is the last release that supports:

- Intel® Ethernet Network Adapter E810-XXV-2 for OCP 2.0

Release 26.4

Release 26.4 is the last release that supports the following:

- SUSE Linux Enterprise Server 11
- Canonical Ubuntu 16.04
- The Microsoft Windows e1q driver and devices. These drivers will no longer be tested or updated. This affects devices based on the following:
    - Intel® 82574L Gigabit Network Connection
    - Intel® Gigabit ET2 Quad Port Server Adapter
    - Intel® 82575EB Gigabit Backplane Connection
    - Intel® 82575EB Gigabit Network Connection
    - Intel® 82576 Gigabit Dual Port Network Connection
    - Intel® 82576 Gigabit Dual Port Server Network Connection
    - Intel® 82576NS Gigabit Ethernet Controller
    - Intel® 82576NS SerDes Gigabit Ethernet Controller
    - Intel® 82583V Gigabit Network Connection
    - Intel® Gigabit CT Desktop Adapter
    - Intel® Gigabit CT2 Desktop Adapter
    - Intel® Gigabit EF Dual Port Server Adapter
    - Intel® Gigabit ET Dual Port Network Connection
    - Intel® Gigabit ET Dual Port Server Adapter
    - Intel® Gigabit ET Quad Port Mezzanine Card

- ◦ Intel® Gigabit ET Quad Port Server Adapter
- ◦ Intel® Gigabit VT Quad Port Server Adapter

- Starting with Release 26.4, the following adapters and devices will no longer be tested or updated. The drivers may still be provided in your download package or on your install media for your convenience.
    - Intel® Ethernet Controller X540-AT2
    - Intel® Ethernet Controller X540-AT1
    - Intel® Ethernet Converged Network Adapter X540-T2
    - Intel® Ethernet Converged Network Adapter X540-T1
    - Intel® X540 Virtual Function

Release 26.3

Release 26.3 is the last release that supports the following:

- Microsoft Windows 10, Version 1803
- Microsoft Windows 10, Version 1903
- Red Hat* Enterprise Linux* (RHEL) 6.x
- The FreeBSD em driver. Maintenance for this driver will be continued by the community.

Release 25.2

Release 25.2 is the last release that supports the following:

- DOS tools and diagnostics
- The 32-bit Microsoft Windows 10 e1d driver. 64-bit Microsoft Windows 10 is still supported. This affects devices based on the following controllers:
    - Intel® Ethernet Connection I217-LM
    - Intel® Ethernet Connection I217-V
    - Intel® Ethernet Connection I218-LM
    - Intel® Ethernet Connection I218-V
    - Intel® Ethernet Connection I219-LM
    - Intel® Ethernet Connection I219-V
- The Linux e1000e driver. Maintenance for this driver will be continued by the community.

Release 25.0

Release 25.0 is the last release to support:

- Microsoft* Windows* 7
- Microsoft* Windows Server* 2008 R2
- Intel® QSFP+ Configuration Utility (QCU)

- IOUtil

Release 24.4

Release 24.4 removed the following from the installation package:

- Support for PRO/100 devices
- DOS drivers
- UEFI driver support for PCI/PCI-X devices
- Support for WinCE
- Microsoft* Windows* 10 RS3 (NDIS65) Universal Drivers. Please use the NDIS68 drivers.
- Support for FCoE

Release 23.5

Starting with Release 23.5, the drivers for the following adapters and devices will no longer be tested or updated. The drivers may still be provided in your download package or on your install media for your convenience.

- Intel® 82566MM Gigabit Network Connection
- Intel® PRO/1000 PB Dual Port Server Connection
- Intel® PRO/1000 EB Network Connection with I/O Acceleration
- Intel® PRO/1000 EB1 Network Connection with I/O Acceleration
- Intel® PRO/1000 EB Backplane Connection with I/O Acceleration
- Intel® 82567V-3 Gigabit Network Connection
- Intel® 82567V-4 Gigabit Network Connection
- Intel® 82567LM-4 Gigabit Network Connection
- Intel® 82567LF Gigabit Network Connection
- Intel® 82567V Gigabit Network Connection
- Intel® 82567LM-2 Gigabit Network Connection
- Intel® 82567LF-2 Gigabit Network Connection
- Intel® 82567V-2 Gigabit Network Connection
- Intel® 82567LM-3 Gigabit Network Connection
- Intel® 82567LF-3 Gigabit Network Connection
- Intel® 82577LM Gigabit Network Connection
- Intel® 82577LC Gigabit Network Connection
- Intel® 82578DM Gigabit Network Connection
- Intel® 82578DC Gigabit Network Connection
- Intel® 82567LM Gigabit Network Connection
- Intel® 82566DM Gigabit Network Connection

- Intel® 82566DC Gigabit Network Connection
- Intel® 82566MC Gigabit Network Connection
- Intel® PRO/1000 PT Dual Port Network Connection
- Intel® PRO/1000 PT Dual Port Server Adapter
- Intel® PRO/1000 PT Dual Port Server Connection
- Intel® PRO/1000 PF Dual Port Server Adapter
- Intel® PRO/1000 PT Network Connection
- Intel® PRO/1000 PT Server Adapter
- Intel® PRO/1000 PF Network Connection
- Intel® PRO/1000 PF Server Adapter
- Intel® PRO/1000 PB Server Connection
- Intel® PRO/1000 PM Network Connection
- Intel® PRO/1000 PL Network Connection
- Intel® PRO/1000 PT Quad Port Server Adapter
- Intel® PRO/1000 PF Quad Port Server Adapter
- Intel® PRO/1000 PT Desktop Adapter
- Intel® PRO/1000 EB1 Backplane Connection with I/O Acceleration
- Intel® PRO/1000 PT Quad Port LP Server Adapter
- Intel® 82566DM-2 Gigabit Network Connection
- Intel® 82562V 10/100 Network Connection
- Intel® 82562V-2 10/100 Network Connection
- Intel® 82562G-2 10/100 Network Connection
- Intel® 82562GT-3 10/100 Network Connection
- Intel® 82562GT-2 10/100 Network Connection
- Intel® 82562GT 10/100 Network Connection
- Intel® 82562G 10/100 Network Connection
- Intel® Gigabit PT Quad Port Server ExpressModule
- Intel® 82566DC-2 Gigabit Network Connection
- Intel® 82598EB 10 Gigabit AF Dual Port Network Connection
- Intel® 10 Gigabit XF SR Dual Port Server Adapter
- Intel® 10 Gigabit XF SR Server Adapter
- Intel® 82598EB 10 Gigabit AF Network Connection
- Intel® 10 Gigabit AT Server Adapter
- Intel® 82598EB 10 Gigabit AT Network Connection
- Intel® 10 Gigabit AT2 Server Adapter
- Intel® 82598EB 10 Gigabit AT2 Network Connection

- Intel® 82598EB 10 Gigabit AT CX4 Network Connection
- Intel® 10 Gigabit SR Dual Port Express Module
- Intel® 10 Gigabit CX4 Dual Port Server Adapter
- Intel® 82598EB 10 Gigabit KX4 Network Connection
- Intel® 10 Gigabit AF DA Dual Port Server Adapter
- Intel® 10 Gigabit XF LR Server Adapter
- Intel® 82598EB 10 Gigabit Dual Port Network Connection

Installation

This section covers how to install Intel® Ethernet adapters, drivers, and other software.

At a high level, installation involves the steps shown below, which are covered in more detail later in other pages.

- If you are installing a network adapter, follow this procedure from step 1.
- If you are upgrading the driver software, start with step 5.

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 appropriate network cable(s).
5. Install the driver.
6. For Windows systems, install the Intel® PROSet software.

Refer to the subsections below for details.

Note:

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

- Hardware Compatibility
- Install the Adapter
    - Select the Correct Slot
    - Insert the Adapter into the Computer
    - PCI Hot Plug Support
        - PCI Hot Plug Support for Microsoft Windows* Operating Systems
- Connect Network Cables
    - Connect Supported SFP+, SFP28, QSFP+, and QSFP28 Modules
    - Connect the Direct Attach Cable

    - Connect the RJ-45 Network Cable
    - Connect the Fiber Optic Network Cable
- Install Drivers and Software
    - Intel Download Center
    - Install via the Complete Driver Pack
    - Install Individual Drivers by OS
        - Install Linux* Drivers
        - Install FreeBSD Drivers
        - Install VMware Drivers
        - Install Windows* Drivers
    - Install Tools and Other Software

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

After completing the basic steps above, you can set up advanced features, if necessary. The available features and the configuration process varies with the device and your operating system. Refer to Device Features for more information; for Linux and FreeBSD drivers, refer to the README file inside the driver tarball.

Hardware Compatibility

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

- The latest BIOS for your system
- One open PCI Express slot (see the specifications of your card for slot compatibility)

Note:

- For devices that support bifurcation, make sure PCI slot bifurcation is enabled in your system's BIOS.
- 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.

Install 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.

   Note:
   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.

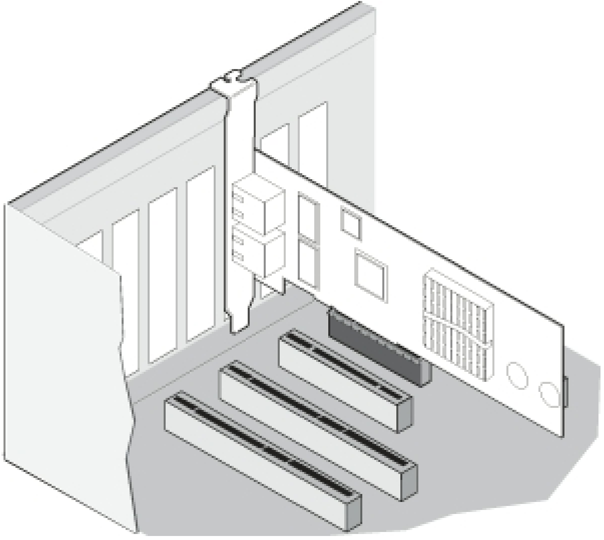

   Note:
   Some PCI Express adapters may have a short connector, making them more fragile than PCI adapters. Excessive force could break the connector. Use caution when pressing the board in the slot.

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.

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 Static Link Aggregation (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® Ethernet 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.

Note:

- 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 adapter is a member of the team before connecting it to the switch.

Connect Network Cables

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

Connect Supported SFP+, SFP28, QSFP+, and QSFP28 Modules

See the Feature Support Matrix for your device family for more information on supported media types.

Connect the Direct Attach Cable

Insert the Direct Attach Cable as shown:

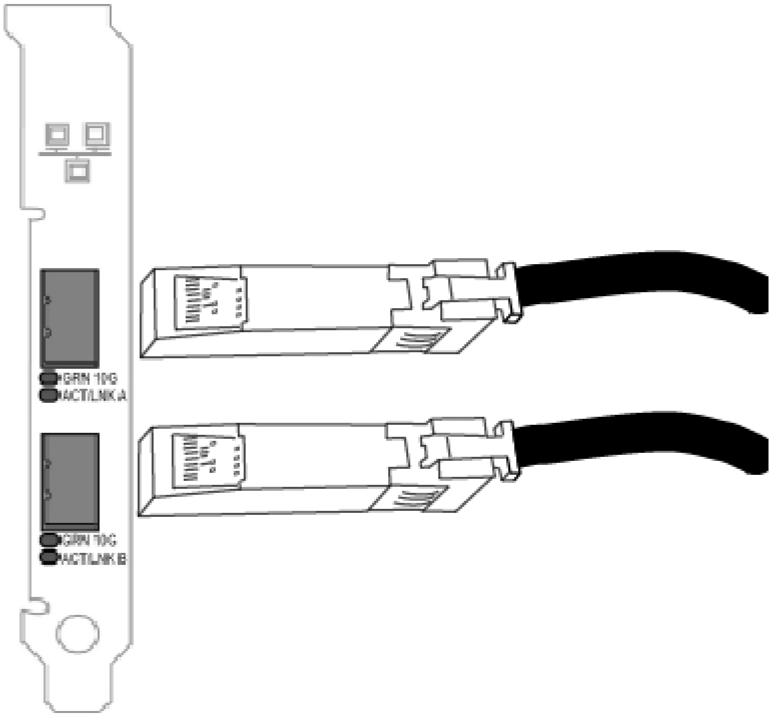

The following table shows the types of direct attached cabling you can use.

Speed	Cable Type	Max Cable Length
100Gbps	QSFP28 Direct Attach Cable	5 meters
40Gbps	SFP+ Direct Attach Cable (Twinaxial)	7 meters
25Gbps	SFP28 Direct Attach Cable (Twinaxial)	5 meters
10Gbps	SFP+ Direct Attach Cable (Twinaxial)	7 meters

Note:

For optimal performance with 25Gbps SFP28 cables, you must use CA-25G-L with RS-FEC and 25GBASE-C.

Connect the RJ-45 Network Cable

Connect the RJ-45 network cable as shown:

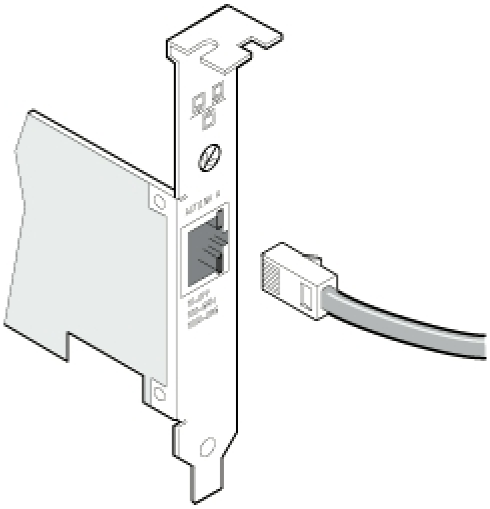

The following table shows the maximum lengths for each cable type at a given transmission speed.

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

Note:

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

Note:

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:

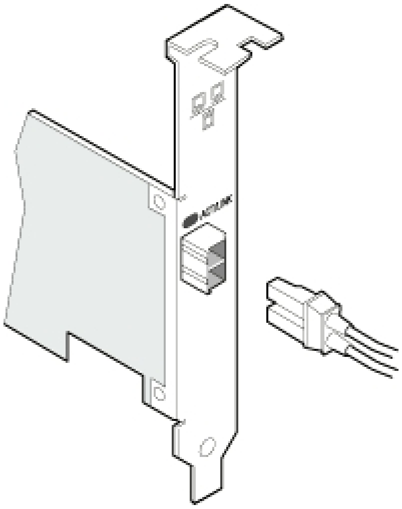

Most connectors and ports are keyed for proper orientation. If the cable you are using is not keyed, check to be sure the connector is oriented properly (transmit port connected to receive port on the link partner, and vice versa).

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.

The following table shows the connection requirements for fiber optic cables.

Device	Laser Wavelength	Connector Type	Cable Type	Max Cable Length
Intel Ethernet LR Server Adapters	1310 nanometer (not visible)	LC	Single-mode fiber with 9.0µm core diameter	10 km
Intel Ethernet 10 Gigabit SR Server Adapters	850 nanometer (not visible)	LC or SC	• Multi-mode fiber with 62.5µm core diameter • Multi-mode fiber with 50µm core diameter	• 33 m • 300 m
Intel Ethernet Gigabit SR Server Adapters	850 nanometer (not visible)	LC or SC	• Multi-mode fiber with 62.5µm core diameter • Multi-mode fiber with 50µm core diameter	• 275 m • 550 m

Install Drivers and Software

There are multiple ways to download drivers and tools for Intel® Ethernet software releases:

- **Download and install the complete driver pack.**
- **Download and install individual drivers.** Exact steps will vary by operating system.

Intel Download Center

You can download all publicly available software from the Intel Download Center at https://www.intel.com/content/www/us/en/download-center/home.html.

Install via the Complete Driver Pack

Downloading the complete driver pack from the Intel Download Center will include drivers for Microsoft Windows*, Linux*, and FreeBSD*, but it is a very large download. We recommend downloading smaller files for your operating system if you don't need software for every OS.

The Intel® Ethernet Adapter Complete Driver Pack is available at https://www.intel.com/content/www/us/en/download/15084/intel-ethernet-adapter-complete-driver-pack.html.

Install Individual Drivers by OS

See the following subsections for details.

- Install Linux* Drivers
  - Locations for Linux Driver Files
  - Manually Build the Linux Driver from Source
  - Build a Binary RPM Package of the Linux Driver
  - Install Linux Drivers via Prebuilt RPM Packages
  - Linux Secure Boot Mode
- Install FreeBSD Drivers
- Install VMware Drivers
- Install Windows* Drivers
  - Before Installing Drivers
  - Install via the Complete Driver Pack
  - Install Only Windows Drivers
  - Install Base Drivers via the Command Line
  - Command Line Syntax for SaveRestore.ps1
  - Examples for SaveRestore.ps1

Install Tools and Other Software

Intel provides a number of tools and applications that allow you to configure or debug the Intel Ethernet devices in your system.

See Tools & Apps for more information about available tools and how to install them.

Tools may be available as individual downloads from the Intel Download Center, or they are included as part of the complete driver pack. Publicly available tools are located in the APPS subfolder of the complete driver pack.

Install Linux* Drivers

The instructions below apply to the following device series and drivers:

Device Series	Driver Name
Intel Adapter Virtual Functions	iavf
Intel Ethernet 800 Series	ice
Intel Ethernet 700 Series	i40e
Intel Ethernet 500 Series	ixgbe
Intel Ethernet 300 Series	igb

You have several options to install Linux* drivers:

- Manually Build the Linux Driver from Source
- Build a Binary RPM Package of the Linux Driver
- Install Linux Drivers via Prebuilt RPM Packages

In the instructions below:

- <driver> is the driver name, such as ice or iavf
- x.x.x is the driver version as indicated in the name of the driver tar file

Locations for Linux Driver Files

Source code for Linux drivers is available to download from the following locations:

- Intel Download Center
- GitHub at https://intel.github.io/ethernet-linux

Note: RPM packages are not available on GitHub.

Manually Build the Linux Driver from Source

1. Move the base driver tar file to the directory of your choice. For example, use /home/username/<driver> or /usr/local/src/<driver>.
2. Untar/unzip the archive:

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

3. Change to the driver src directory:

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

4. Compile the driver module:

   ```
   make install
   ```

   The binary will be installed as:

   ```
   /lib/modules/<KERNEL VER>/updates/drivers/net/ethernet/intel/<driver>/<driver>.ko
   ```

   The install location listed above is the default location. This may differ for various Linux distributions.

   Note: For ice devices:

   - To build the driver using the schema for unified ethtool statistics, use the following command:

     ```
     make CFLAGS_EXTRA='-DUNIFIED_STATS' install
     ```

   - To compile the ice driver with ADQ (Application Device Queues) flags set, use the following command, where `<nproc>` is the number of logical cores:

     ```
     make -j<nproc> CFLAGS_EXTRA='-DADQ_PERF_COUNTERS' install
     ```

     (This will also apply the above `make install` command.)

   Note:

   For i40e devices: To gather and display additional statistics, use the `I40E_ADD_PROBES` pre-processor macro:

   ```
   make CFLAGS_EXTRA=-DI40E_ADD_PROBES
   ```

   Please note that this additional statistics gathering can impact performance.

   Note:

   For devices that support RDMA: You may see warnings from **depmod** related to unknown RDMA symbols during the make of the out-of-tree (OOT) base driver. These warnings are normal and appear because the in-tree RDMA driver will not work with the OOT base driver. To address the issue, you need to install the latest OOT versions of the base and RDMA drivers.

5. Load the module using the **modprobe** command.

   To check the version of the driver and then load it:

   ```
   modinfo <driver>modprobe <driver>
   ```

   Alternately, make sure that any older device drivers are removed from the kernel before loading the new module:

   ```
   rmmod <driver>; modprobe <driver>
   ```

6. Assign an IP address to the interface by entering the following, where <ethX> is the interface name that was shown in dmesg after modprobe:

   ```
   ip address add <IP_address>/<netmask bits> dev <ethX>
   ```

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

   ```
   ping <IP_address>
   ```

Note:

For certain distributions like (but not limited to) Red Hat* Enterprise Linux* 7 and Ubuntu*, once the driver is installed, you may need to update the initrd/initramfs file to prevent the OS loading old versions of the driver.

For Red Hat distributions:

```
dracut --force
```

For Ubuntu:

```
update-initramfs -u
```

Build a Binary RPM Package of the Linux Driver

Note: RPM functionality has only been tested in Red Hat distributions.

1. Run the following command, where <x.x.x> is the version number for the driver tar file:

```
rpmbuild -tb <driver>-<x.x.x>.tar.gz
```

Note:

For the build to work properly, the currently running kernel MUST match the version and configuration of the installed kernel sources. If you have just recompiled the kernel, reboot the system before building.

2. After building the RPM, the last few lines of the tool output contain the location of the RPM file that was built. Install the RPM with one of the following commands, where <RPM> is the location of the RPM file:

```
rpm -Uvh <RPM>
```

or:

```
dnf/yum localinstall <RPM>
```

3. **For ice, i40e, and iavf devices only**: If your distribution or kernel does not contain inbox support for auxiliary bus, you must also install the auxiliary RPM:

```
rpm -Uvh <driver RPM> <auxiliary RPM>
```

or:

```
dnf/yum localinstall <driver RPM> <auxiliary RPM>
```

Note:

On some distributions, the auxiliary RPM may fail to install due to missing kernel-devel headers. To workaround this issue, specify --excludepath during installation. For example:

```
rpm -Uvh auxiliary-1.0.0-1.x86_64.rpm--excludepath=/lib/modules/3.10.0-957.el7.x86_64/source/include/linux/auxiliary_bus.h
```

- To compile the driver on some kernel/arch combinations, you may need to install a package with the development version of **libelf** (e.g., **libelf-dev**, **libelf-devel**, **elfutils-libelf-devel**).
- When compiling an out-of-tree driver, details will vary by distribution. However, you will usually need a kernel-devel RPM or some RPM that provides the kernel headers at a minimum. The RPM kernel-devel will usually fill in the link at /lib/modules/'uname -r'/build.

Install Linux Drivers via Prebuilt RPM Packages

Intel provides prebuilt RPM Package Manager (RPM) files for the following distributions on the Intel Download Center:

- For Red Hat distributions: KMOD RPM files
- For SUSE distributions: Kernel Module Package (KMP) RPM files

See the following webpages on the RHEL and SLES websites for more information on installing RPM packages:

- Red Hat: How to install or upgrade an RPM package?
- SUSE: Introduction to RPM Packaging

Linux Secure Boot Mode

This software release includes RPM packages that contain:

- Device driver signed with Intel's private key in precompiled kernel module form
- RDMA driver
- Complete source code for above drivers
- Intel's public key

This release includes the Intel public key to allow you to authenticate the signed driver in secure boot mode. To authenticate the signed driver, you must place Intel's public key in the UEFI Secure Boot key database.

- The driver kernel module for a specific kernel version can be used with errata kernels within the same minor OS version, unless the errata kernel broke kABI. Whenever you update your kernel with an errata kernel, you must reinstall the driver RPM package.
- The RDMA driver will be installed if you reinstall the driver RPM package. If you want to remove the RDMA driver, you will have to do so every time you install the RPM package (for example, when you update your kernel with an errata kernel).
- If you decide to recompile the *.ko* module from the provided source files, the new *.ko* module will not be signed with any key. To use this *.ko* module in Secure Boot mode, you must sign it yourself with your own private key and add your public key to the UEFI Secure Boot key database.

Install FreeBSD Drivers

Note:

The FreeBSD driver package is to be used only as a standalone archive and the user should not attempt to incorporate it into the kernel source tree.

The instructions below apply to the following device series and drivers:

Device Series	Driver Name
Intel Adapter Virtual Functions	iavf
Intel Ethernet 800 Series	ice
Intel Ethernet 700 Series	ixl
Intel Ethernet 500 Series	ix
Intel Ethernet 300 Series	igb

In the instructions below:

- <driver> is the driver name, such as ice or iavf
- x.x.x is the driver version as indicated in the name of the driver tar file

1. Move the base driver tar file to the directory of your choice. For example, use /home/username/<driver> or /usr/local/src/<driver>.
2. Untar/unzip the archive:

   ```
   tar xzf <driver>-x.x.x.tar.gz
   ```

   This will create the <driver>-x.x.x directory.
3. To install the man page:

   ```
   cd <driver>-x.x.xgzip -c <driver>.4 > /usr/share/man/man4/<driver>.4.gz
   ```

4. To load the driver onto a running system:

   ```
   cd <driver>-x.x.xmakekldload ./if_<driver>.ko
   ```

   For the iavf driver, to install the driver without using iflib:

   ```
   cd iavf-x.x.x/srcmake legacykldload ./if_iavf.ko
   ```

Note:

Ensure the driver isn't compiled into the currently running kernel. You can do that by adding nodevice <driver> to your kernel configuration file and rebuilding your kernel. See sections 10.4 and 10.5 in Configuring the FreeBSD Kernel in the FreeBSD kernel documentation for more information.

Note:

For ice devices, running the **make** command will not install the Dynamic Device Personalization (DDP) package and could cause the driver to fail to load. See step #7 below for more information.

5. To assign an IP address to the interface, enter the following, where X is the interface number for the device:

```
ifconfig <driver>X <IP_address>
```

For example:

```
ifconfig ice0 <IP_address>
```

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

```
ping <IP_address>
```

7. If you want the driver to load automatically when the system is booted, do the following.

   a. For all devices except ice:

   ```
   cd <driver>-x.x.x/srcmakemake install
   ```

   For ice devices:

   ```
   cd <driver>-x.x.xmakemake install
   ```

Note:

For ice devices, it's important to do make install so that the driver loads the DDP package automatically.

   b. Edit /boot/loader.conf and add the following line:

   ```
   if_<driver>_load="YES"
   ```

8. If you want the device to connect to a network without manual intervention after a boot, edit /etc/rc.conf and create the appropriate ifconfig_<driver>X entry:

```
ifconfig_<driver>X="<ifconfig_settings>"
```

For example:

```
ifconfig_ice0="inet 192.168.10.1 netmask 255.255.255.0"
```

Note: For assistance, see the ifconfig man page.

Install VMware Drivers

ESXi drivers for Intel Ethernet devices are only available from VMware's download site: https://customerconnect.vmware.com/home.

Please refer to VMware's download site for the latest drivers and installation information.

Install Windows* Drivers

There are multiple ways to install device drivers on Microsoft Windows:

- Install via the Complete Driver Pack. This option will install Windows* drivers and Intel® PROSet but is a very large download.
- Install Only Windows Drivers. This option will not install Intel PROSet but is a much smaller download.
- Install Base Drivers via the Command Line
- You can also Save and Restore an Adapter's Configuration Settings via the command line.

This page describes the above installation methods for Windows drivers. See Installing Intel® PROSet for instructions on how to install Intel PROSet.

Note:

- To successfully install or uninstall the drivers or software, you must have administrative privileges on the computer completing installation.
- Installing the drivers will update the drivers for all supported Intel Ethernet 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 is present on the system. Before you use the Roll Back Driver feature, remove any teams. Then remove Intel PROSet using **Programs and Features** from the Control Panel of Windows. See Installing Intel® PROSet for details regarding Intel PROSet.
- 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 Drivers

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 procedures below to install the driver. Drivers for all Intel adapters supported by this software release are installed.

Install via the Complete Driver Pack

To download and install via the complete driver pack:

1. Download the latest software package from the support website and transfer it to the system. See Install via the Complete Driver Pack for the URL for the complete driver pack.
2. Extract the downloaded software package to your hard drive.
3. If the Found New Hardware Wizard screen is displayed, click **Cancel**.
4. Navigate in your extracted files to \APPS\SETUP\SETUPBD and then the Windows subfolder corresponding to your version of Windows (32-bit or 64-bit).
5. Inside the Winx64 or Win32 folder, double-click on *SetupBD.exe*.
6. Complete the steps in the installation wizard.
7. If you want to install Intel PROSet, navigate to \APPS\PROSETDX and then the Windows subfolder corresponding to your version of Windows (32-bit or 64-bit).
   - See Ethernet Cmdlets for Intel® Ethernet and Intel® PROSet for more information about this tool.
8. Inside the Winx64 or Win32 folder, double-click on *DxSetup.exe*.
9. Complete the steps in the installation wizard.

Install Only Windows Drivers

To download and install only Windows drivers:

1. Go to the Intel Download Center and find the relevant download for your version of Windows. See Intel Download Center for the link.
2. Download and extract the **Wired_driver_XX.X_*.zip** file for your version of Windows, where XX.X is the release number. This file will install the base driver(s) for your system.
   - If you are running a 32-bit operating system, download **Wired_driver_XX.X_32.zip**.
   - If you are running a 64-bit operating system, download **Wired_driver_XX.X_x64.zip**.
3. In the extracted driver files, double-click on the **.exe** file to launch the installation.

4. In the dialog box that opens, click on **OK** to install the drivers.
5. The device driver(s) will install. Click **Close** when prompted.

> **Note:**
>
> This method does not install Intel PROSet. See Installing Intel® PROSet for additional instructions.

Install Base Drivers via the Command Line

SetupBD.exe is a file that allows you to install the Windows base drivers from the command line. SetupBD is only available when you download and Install via the Complete Driver Pack.

To install Windows base drivers using SetupBD:

1. Download the complete driver pack. See Install via the Complete Driver Pack for details.
2. Extract the downloaded files to your desired location.
3. In the downloaded files, navigate to APPS\SETUP\SETUPBD and the appropriate subfolder for your version of Windows (32-bit or 64-bit).
4. From the command line, call SetupBD using the following syntax:

   ```
   SetupBD [parameters]
   ```

   See below for supported command line parameters and examples.

Parameters

/h, /?

Displays help file for SetupBD.

/l <*path\filename*>

Create a log file with the specified path and filename. If you do not specify the path and filename, SetupBD creates a log file (named *SetupBD_<timestamp>.log*; for example, *SetupBD_18-04-2022_14-29-20.log*) in the current directory.

/m

Non-interactive install. This still displays the installation GUI, but you cannot interact with it. Use the /s switch to suppress the GUI.

/n

Ignore INF excludes and force scan all *.inf* files.

/nr

Suppress the system reboot. If you include this switch you will need to manually reboot the system for the changes to take effect. Must be used with the /s switch. This switch is ignored if it is included with the /r switch.

/r

Force a system reboot after the installation completes. Must be used with the /s switch.

/s

Silent installation. Suppresses all installation messages and errors.

Examples

To install and/or update the Windows driver(s) and display the GUI:

```
SetupBD
```

To install and/or update the Windows driver(s) silently:

```
SetupBD /s
```

To install and/or update the Windows driver(s) silently and create a log file in `c:\temp:

```
SetupBD /s /l c:\temp\install.log
```

To install and/or update the Windows driver(s) silently and force a reboot:

```
SetupBD /s /r
```

To install and/or update the Windows driver(s) silently and force a reboot (/nr is ignored):

```
SetupBD /s /r /nr
```

Save and Restore an Adapter's Configuration Settings

The Save and Restore Command Line Tool (*SaveRestore.ps1*) 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. A saved configuration file can be used to restore adapter settings after an operating system upgrade. However, all adapter configuration settings may not be restored depending on the features supported by the new operating system or adapter configuration software.

Note:

- 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 Intel 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.

- Intel PROSet must be installed for the *SaveRestore.ps1* script to run. See Installing Intel® PROSet for installation instructions.
- 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 for SaveRestore.ps1

Use the following syntax to call *SaveRestore.ps1*:

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

Parameters

-Action

Required. Valid values are:

save:

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.

restore:

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 an adapter to a different slot, this may result in the script applying the saved settings to a different device.

Note:

- 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 for SaveRestore.ps1

Save Example

To save the adapter settings to a file on a removable media device:

1. Open a Windows PowerShell Prompt.
2. Navigate to the directory where *SaveRestore.ps1* is located (generally c:\Program Files\Intel\Wired Networking\PROSET).
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:

1. Open a Windows PowerShell Prompt.
2. Navigate to the directory where *SaveRestore.ps1* is located (generally c:\Program Files\Intel\Wired Networking\PROSET).
3. Type the following:

```
SaveRestore.ps1 -Action Restore -ConfigPath e:\settings.txt
```

Release Notes

View the Intel® Ethernet Controller Products Release Notes on intel.com for information on added or removed features, supported operating systems for each driver, bug fixes, and known issues and limitations for the software release.

You can view the latest release or past releases. To view the release notes for past releases:

1. Go to the following page on intel.com:
   https://www.intel.com/content/www/us/en/search.html?ws=idsa-default#q=Intel® Ethernet Controller Products Release Notes&sort=relevancy&f:@tabfilter=[Developers]
2. In the list of results, click on the link for “Intel® Ethernet Controller Products Release Notes.”
   - By default, the page lists the newest released version.
   - If you need to view the release notes for a different version, click on the “More Versions” link in the right side of the search result and select your desired version.

Known Issues and Limitations

See the Intel® Ethernet Controller Products Release Notes for all known issues and fixed issues.

Feature Support Matrix

We provide Feature Support Matrix documents for many of our product families. These documents detail supported features, cables, media types, operating systems, and other technical information.

View the following documents in the Intel Resource and Documentation Center:

- Intel® Ethernet 800 Series:
  - Intel® Ethernet Controller E810 Feature Support Matrix
  - Intel® Ethernet Connection E82X Feature Support Matrix
- Intel® Ethernet 700 Series:
  - Intel® Ethernet Controller X710/XXV710/XL710 Feature Support Matrix
  - Intel® Ethernet Connection X722 Feature Support Matrix
  - Intel® Ethernet Controller X710-TM4/AT2 and V710-AT2 Feature Support Matrix
- Intel® Ethernet 500 Series:
  - Intel® Ethernet Controller X550 Feature Support Matrix

Device Features

This section describes the features available on Intel® Ethernet devices. Major features are organized alphabetically.

Note:

- Available settings are dependent on your device and operating system. **Not all settings are available on every device/OS combination.**
- Some features in this section refer to Intel® PROSet, Intel® PROSet for Windows* Device Manager, Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU), Intel® PROSet for Windows PowerShell* software, or Ethernet Cmdlets for Intel® Ethernet. Refer to Ethernet Cmdlets for Intel® Ethernet and Intel® PROSet for more information on each of these applications.

- Adapter Teaming
  - Teaming Modes
  - Configuring Teams with Intel® PROSet
- Adaptive Inter-Frame Spacing
- Data Center Bridging (DCB)
- Direct Memory Access (DMA) Coalescing
- Dynamic Device Personalization (DDP)
- Firmware
  - Firmware Security
  - Firmware Rollback Mode
  - Firmware Recovery Mode
  - Using Devlink to Update a Device's NVM

-
    - Flash Images
    - Intel® Boot Agent
- Remote Direct Memory Access (RDMA)
- Accessing Remote NVM Express* Drives Using RDMA
- Setting Speed and Duplex
- Thermal Monitoring
- Timestamps
    - PTP Hardware Timestamp
    - Software Timestamp
- Transmit Buffers
- UEFI Network Device Drivers
- VF Loopback Pacing
- Virtualization Support
    - Single Root I/O Virtualization (SR-IOV)
    - Virtual Machine Queue Offloading
    - Using Intel Network Adapters in a Microsoft* Hyper-V* Environment
- Virtual LANs (VLANs)
- Wait for Link

Adapter Teaming

Intel® Advanced Network Services (Intel® ANS) teaming lets you take advantage of multiple adapters in a system by grouping them together. Intel 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.

For more information on VLANs, see Virtual LANs (VLANs).

See the following subsections for details.

- Teaming Modes
    - Overview
    - Adapter Fault Tolerance (AFT)
    - Switch Fault Tolerance (SFT)
    - Adaptive/Receive Load Balancing (ALB/RLB)
    - Virtual Machine Load Balancing (VMLB)
    - Static Link Aggregation (SLA)
    - IEEE 802.3ad Dynamic Link Aggregation

- Multi-Vendor Teaming (MVT)
- Configuring Teams with Intel® PROSet
  - Configuring Teams with Intel® PROSet for Windows* Device Manager
  - Configuring Teams with Windows PowerShell*
  - Configuring Teams with Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU)

Configuration and Compatibility Notes

- Microsoft Windows* 10 is the last Windows operating system version that supports Intel ANS. Intel ANS is not supported on Microsoft Windows 11 and later.
- Intel ANS is not supported 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. Intel adapters that do not support Intel ANS 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.
- 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 the feature 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. When you physically remove an adapter that is part of a team or a VLAN, you must reboot or reload the team/VLAN before using that adapter in the same network. This will prevent Ethernet address conflicts.
- 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.
- Data Center Bridging will be automatically disabled when an adapter is added to a team with non-DCB capable adapters.
- 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.
- If you want to set up VLANs on a team, you must first create the team.
- After adding a VLAN to the team, the Network Connections window shows the team as disabled or network cable unplugged. This is normal. The connection protocols are now bound to the VLAN on the team. You can configure the connection protocols in the Properties for the VLAN.

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. If the adapter is part of a team, remove the adapter from the team.
2. Shut down the system and unplug the power cable.
3. Disconnect the network cable from the adapter.
4. Open the case and remove the adapter.
5. Insert the replacement adapter. (Use the same slot, otherwise Windows assumes that there is a new adapter.)
6. Reconnect the network cable.
7. Close the case, reattach the power cable, and power up the server.

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 Intel 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 Intel ANS teams and VLANs involved in the LBFO team and recreate them. This is an optional (all bindings are restored when the LBFO team is removed), but strongly recommended step

Note:

- If you add an Intel AMT enabled port to an LBFO team, do not set the port to Standby in the LBFO team. If you set the port to Standby you may lose AMT functionality.
- DCB is incompatible with Microsoft Server LBFO Teams. Do not create an LBFO team when DCB is installed. Do not install DCB if you use LBFO teaming. Install failures and persistent link loss may occur if DCB and LBFO are used on the same port.

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.
- TOE (TCP Offload Engine) enabled devices cannot be added to an Intel 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.
- 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

Overview

The following teaming modes are supported, and are described later in this page:

- Adapter Fault Tolerance (AFT)
- Switch Fault Tolerance (SFT)
- Adaptive/Receive Load Balancing (ALB/RLB)
- Virtual Machine Load Balancing (VMLB)
- Static Link Aggregation (SLA)
- IEEE 802.3ad Dynamic Link Aggregation

- Multi-Vendor Teaming (MVT)

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 offloads. These settings are available in the Advanced tab in Intel® PROSet. *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.

Note:

- 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. Under normal conditions, the Primary adapter handles all traffic. The Secondary adapter will receive 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.

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.

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

Note:

- 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 adapter 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 adapters 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.

Note:

- 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 (VMLB)

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.

Note:

- 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 an Intel AMT enabled adapter in a VMLB team.

Static Link Aggregation (SLA)

Static Link Aggregation (SLA) is very similar to Adaptive/Receive Load Balancing (ALB/RLB), 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

Note:

- 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 1Gbps and 10Gbps, but two aggregators will be formed, one for each speed. Also, if a team contains 1Gbps ports connected to one switch, and a combination of 1Gbps and 10Gbps ports connected to a second switch, three aggregators would be formed. One containing all the ports connected to the first switch, one containing the 1Gbps ports connected to the second switch, and the third containing the 10Gbps ports connected to the second switch.

Note:

- 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 (MVT)

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.

Configuring Teams with Intel® PROSet

This page describes how to configure Intel® Advanced Network Services (Intel® ANS) teams using Intel® PROSet.

Refer to the following subsections for more specific information:

- Configuring Teams with Intel® PROSet for Windows* Device Manager
- Configuring Teams with Windows PowerShell*
- Configuring Teams with Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU)

Configuring Teams with Intel® PROSet for Windows* Device Manager

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.

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.

Changing Which Adapters Are in a Team

1. Launch Windows Device Manager.

2. Open the Team Properties dialog box by double-clicking on a team listing in the Computer Management window.
3. Click the **Settings** tab.
4. Click **Modify Team**, then click the **Adapters** tab.
5. 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.
6. Click **OK**.

Renaming a Team

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

Note:

If you modify a team name from the team property sheet, it may take several minutes for the name to change in Device Manager. Closing and opening Device Manager will load the new 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.

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 standalone mode.

Specifying a Preferred Primary or Secondary Adapter

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

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.

Configuring Teams with Windows PowerShell*

Using non-Intel cmdlets, such as the Set-NetAdapterAdvancedProperty cmdlet provided in Microsoft PowerShell*, to change settings for an Intel ANS-teamed adapter 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 Intel PROSet Teaming GUI. You can repair the issue by changing the setting back to its previous state, or by removing the adapter from the Intel ANS team and then adding it back.

Creating a Team

Use the New-IntelNetTeam cmdlet. For example:

```
New-IntelNetTeam -TeamMemberNames "<adapter1_name>", "<adapter2_name>"
-TeamMode AdapterFaultTolerance -TeamName "<team_name>"
```

Changing Which Adapters Are in a Team

Use the Add-IntelNetTeamMember or Remove-IntelNetTeamMember cmdlet. For example:

```
Add-IntelNetTeamMember -TeamName "<team_name>" -Name "<adapter_name>"
```

Renaming a Team

Use the Set-IntelNetTeam cmdlet. For example:

```
Set-IntelNetTeam -TeamName "<team_name>" -NewTeamName "<new_team_nam
e>"
```

Removing a Team

Use the Remove-IntelNetTeam 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 standalone mode.

Specifying a Preferred Primary or Secondary Adapter

Use the Set-IntelNetTeam cmdlet. For example:

```
Set-IntelNetTeam -TeamName "Team 1" -PrimaryAdapterName "<adapter1_name>
"-SecondaryAdapterName "<adapter2_name>"
```

Configuring Teams with Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU)

Creating a Team with Intel PROSet ACU

1. Launch Intel PROSet ACU.
2. Select an adapter to start the team.
3. Click the **Teaming/VLANs** tab.
4. In the Teaming panel, click **Create Team**.
5. Select the adapters to include in the team, then click **Next**.
6. Type a name for the team.
7. Select the teaming mode, then click **Next**.
8. [Optional] Designate Primary and Secondary adapters for the team.
9. Click **Finish**.

Changing Which Adapters Are in a Team

1. Launch Intel PROSet ACU.
2. Select the team you wish to modify.
3. In the Team Members panel, click **Modify Members**.
4. Select the adapters that will be members of the team.
5. Click **Apply Changes**.

Renaming a Team

1. Launch Intel PROSet ACU.
2. Select the team you wish to modify.
3. In the Team Information panel, type a new team name.
4. Click **Apply Changes**.

Removing a Team

1. Launch Intel PROSet ACU.
2. Select the team you wish to remove.
3. Click **Remove Team**.

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 standalone mode.

Specifying a Preferred Primary or Secondary Adapter

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

1. Launch Intel PROSet ACU.
2. Select the team you wish to modify.
3. Select your preferred Primary Adapter.
4. Select your preferred Secondary Adapter.
5. Click **Apply Changes**.

Adaptive Inter-Frame Spacing

This setting 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Possible values for this setting are:

- Enabled
- Disabled

Data Center Bridging (DCB)

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).

DCB 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 (LLDP) standard (IEEE 802.1AB) that enable Data Center Bridging Capability Exchange Protocol (DCBX)

Adapter firmware implements LLDP and DCBX protocol agents as per 802.1AB and 802.1Qaz respectively.

There are two supported versions of DCBX:

- CEE Version
- IEEE Version

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

Configuring DCB for Windows*

To change this setting in Intel® PROSet:

This setting is found on the Data Center tab of the device's Device Manager property sheet or in the Data Center panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

You can use Intel PROSet to perform the following tasks:

- Display status:
    - Enhanced Transmission Selection
    - Priority Flow Control
    - 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.

      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

Note:

- On X710 based devices running Microsoft Windows, DCB is only supported on NVM version 4.52 and newer. Older NVM versions must be updated before the adapter is capable of DCB support in Windows.
- If *QOS/DCB is not available, it may be for one of the following reasons:
    - The Firmware LLDP (FW-LLDP) agent was disabled from a pre-boot environment (typically UEFI).
    - This device is based on the Intel® Ethernet Controller X710 and the current link speed is 2.5Gbps or 5Gbps.

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*

Intel Ethernet drivers support firmware-based or software-based DCBX in Linux, depending on the underlying PF device. The following table summarizes DCBX support by driver.

Linux Driver	Firmware-Based DCBX	Software-Based DCBX
ice	Supported	Supported
i40e	Supported	Supported
ixgbe	Not supported	Supported

In **firmware-based** mode, firmware intercepts all LLDP traffic and handles DCBX negotiation transparently for the user. In this mode, the adapter operates in “willing” DCBX mode, receiving DCB settings from the link partner (typically a switch). The local user can only query the negotiated DCB configuration.

In **software-based** mode, LLDP traffic is forwarded to the network stack and user space, where a software agent can handle it. In this mode, the adapter can operate in either “willing” or “nonwilling” DCBX mode and DCB configuration can be both queried and set locally. Software-based mode requires the FW-based LLDP Agent to be disabled and kernel CONFIG_DCB enabled.

Note:

- Only one LLDP/DCBX agent can be active on a single interface at a time.
- Software-based and firmware-based DCBX modes are mutually exclusive.
- When the firmware DCBX agent is active, software agents will not be able to receive or transmit LLDP frames. See Firmware Link Layer Discovery Protocol (FW-LLDP), as well as the Linux driver readme in your installation, for information on enabling or disabling the FW-LLDP agent.
- In software-based DCBX mode, you can configure DCB parameters using software LLDP/DCBX agents that interface with the Linux kernel’s DCB Netlink API. We recommend using OpenLLDP as the DCBX agent when running in software mode. For more information, see the OpenLLDP man pages and https://github.com/intel/openlldp.
- For information on configuring DCBX parameters on a switch, please consult the switch manufacturer’s documentation.

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.

Configuring iSCSI Over DCB in Windows

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.

To change this setting in Intel PROSet:

This setting is found on the Data Center tab of the device’s Device Manager property sheet or in the Data Center panel in Intel PROSet ACU.

This setting provides feedback as to the DCB state, operational or non-operational, as well as providing additional details should it be non-operational.

Note:

On Microsoft Windows Server, if you configure Priority using IEEE, the iSCSI policy may not be created automatically. To create the iSCSI policy manually, use Powershell* and type:

```
New-NetQosPolicy -Name "UP4" -PriorityValue 8021 Action 4 -iSCSI
```

Using iSCSI over DCB with Intel 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 Intel 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.

See Adapter Teaming for more information on Intel ANS teams.

Configuring iSCSI Over DCB in Linux

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-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. 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.

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 Ethernet device 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.

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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "DMA Coalescing"-DisplayValue "Enabled"
```

Dynamic Device Personalization (DDP)

Devices based on the Intel® Ethernet 800 Series require a Dynamic Device Personalization (DDP) package file to enable advanced features (such as dynamic tunneling, Intel Ethernet Flow Director, RSS, and ADQ).

DDP allows you to change the packet processing pipeline of a device by applying a profile package to the device at runtime. Profiles can be used to, for example, add support for new protocols, change existing protocols, or change default settings. DDP profiles can also be rolled back without rebooting the system.

The DDP package loads during device initialization or driver runtime, depending on the operating system. The driver checks to see if the DDP package is present and compatible. If this file exists, the driver will load it into the device. If not, the driver will go into Safe Mode where it will use the configuration contained in the device's NVM.

Safe Mode disables advanced and performance features, and supports only basic traffic and minimal functionality, such as updating the NVM or downloading a new driver or DDP package. For more information, see Safe Mode.

A general-purpose, default DDP package is automatically installed with all supported Intel Ethernet 800 Series drivers on supported operating systems. Additional DDP packages are available to address needs for specific market segments or targeted solutions.

Refer to the following for more information on configuring DDP:

- Intel® Ethernet Controller E810 Dynamic Device Personalization (DDP) Technology Guide
    - Search for "DDP" on the Intel Resource and Documentation Center for additional DDP technology and configuration guides.
- The readme file inside the DDP package zip file

Note:

- If you are using DPDK, see the DPDK documentation at https://www.dpdk.org/ for installation instructions and more information.
- In ESXi:

- Support for DDP packages for specific market segments requires the following:
    - □ Driver: icen 1.9.1.x or higher
    - □ Tool: intnet 1.8.3.x or higher
- Use esxcli to load and unload DDP packages for specific market segments during driver runtime.
- A package update is not persistent between device resets or system reboots.

Firmware

Firmware is a layer of software that is programmed into a device's memory. It provides low level functionality for the device. In most cases you will not notice the firmware on your device at all. Firmware error states usually occur because of an unsuccessful update.

See the following subsections for details.

- Firmware Security
- Firmware Rollback Mode
- Firmware Recovery Mode
- Using Devlink to Update a Device's NVM
- Firmware Logging
- Debug Dump

Firmware Security

Intel or your equipment manufacturer will occasionally release a firmware security patch. We recommend that you update your firmware to the latest version available for your device to take advantage of these security patches. Firmware updates for Intel Ethernet devices will have a Security Revision number (SRev).

Minimum Security Revision Enforcement

Firmware security updates can be undone if you install a previous version of the firmware onto your device. Intel firmware releases include a Minimum Security Revision (MinSRev) enforcement feature. This means you can block someone from installing a lower revision of the firmware onto your device. This will limit the rollback capabilities of your device. The firmware update process will block the update if the supplied firmware has a lower security revision (SRev) than the MinSRev value of the firmware currently loaded on the device. Only update the MinSRev value if you are certain you will not need to roll the firmware back to an earlier version.

You can update the MinSRev value during the firmware update process, locking the current security version in as the new MinSRev baseline, by using the -optinminsrev command line option.

> **Important:**
>
> The MinSRev value on a device can never be decreased. Once the MinSRev is increased, NVM downgrades attempting to install a lower Security revision (SRev) than the current MinSRev will be rejected by the device. Users who want to downgrade firmware without regard to security revisions should not use this feature.

SRev and MinSRev Examples

To view your device's current SRev and MinSRev:

You can use the nvmupdate tool's inventory mode to view your device's current SRev and MinSRev values as follows:

- Windows:

```
nvmupdatew64e -i  -l update.log -o results.xml -c nvmupdate.cfg -optinminsrev
```

- Linux:

```
nvmupdate64e -i  -l update.log -o results.xml -c nvmupdate.cfg
```

Where:

-i

Sets nvmupdate to inventory mode.

-l update.log

Specifies the name of the log file.

-o results.xml

Specifies the name of the results file. This is an XML file that contains the inventory/update results.

-c nvmupdate.cfg

Specifies the name of the configuration file. This is a text file that contains descriptions of networking devices and firmware versions for those devices.

-optinminsrev

Specifies that the MinSRev and SRev values are included in the *results.xml* file.

Examine the *results.xml* file for the SRev and MinSRev values.

> **Note:**
>
> Make sure you specify -i for inventory mode. If you specify -u, the tool will update the MinSRev value, rather than simply disclose it. You can achieve the same results by specifying MINSREV:TRUE in the configuration file.

See Intel® Ethernet NVM Update Tool for more information on how to use the nvmupdate tool.

To update your device's MinSRev:

1. Download and extract the NVM Update Package for your device.
2. Use the command line to update your device's MinSRev:
   - Windows:

     ```
     nvmupdatew64e -u -optinminsrev -l update.log -o results.xml -c nvmupdate.cfg
     ```

   - Linux:

     ```
     nvmupdate64e -u -optinminsrev -l update.log -o results.xml -c nvmupdate.cfg
     ```

Where:

-u

Sets nvmupdate to update mode.

-optinminsrev

Tells the tool to update the MinSRev value.

-l update.log

Specifies the name of the log file.

-o results.xml

Specifies the name of the results file. This is an XML file that contains the inventory/update results.

-c nvmupdate.cfg

Specifies the name of the configuration file. This is a text file that contains descriptions of networking devices and firmware versions for those devices.

See Intel® Ethernet NVM Update Tool for more information on how to use the nvmupdate tool.

Firmware Rollback Mode

When a port is in firmware rollback mode, it may have reduced functionality.

Usually a device enters firmware rollback mode when a firmware update does not complete correctly. Rebooting or power cycling the system may allow the port to use the previous firmware.

You may need to reapply the firmware update to regain full functionality on the device. Use the appropriate NVM Update Package to update the device's firmware. Download the latest NVM Update Package from your vendor's support website and follow the instructions in it.

After restoring the NVM image, you may need to perform an A/C power cycle of the system.

Firmware Recovery Mode

A device will enter Firmware Recovery mode if it detects a problem that requires the firmware to be reprogrammed. 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.

Detecting Recovery Mode

During initialization, a device can enter recovery mode if the device firmware detects a problem with the LAN device, mandating NVM reprogramming to restore normal operation. After thorough internal testing of the NVM (typically less than 10 minutes, but in some cases longer), the device enters Recovery Mode.

Firmware Recovery Mode Errors and Messages

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

"Firmware recovery mode detected. Limiting functionality. Refer to the Intel® Ethernet Adapters and Devices User Guide for details on firmware recovery mode."

"Firmware recovery mode detected. The underlying hardware has been deactivated. Refer to the Intel® Ethernet Adapters and Devices User Guide for details on firmware recovery mode."

"Firmware recovery mode detected. Initialization failed."

"Firmware recovery mode detected. Limiting functionality."

"Initialization failure due to repeated FW resets."

The last message above is usually an indication that the device is about to enter Recovery Mode. The device may be able to return to normal functionality without intervention. This may take several minutes. No action is required unless the device does enter Recovery Mode.

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.

The process for resolution of Firmware Recovery Mode Issues is outlined in the subsections below.

NVM Self Check

The process begins after power-on or reboot. At this time, the firmware will perform tests to assess whether there is damage or corruption of the device N VM image.

Actions:

- If NVM image damage or corruption **is not** detected, the device will initialize and operate normally. No further action is required.
- If NVM image damage or corruption **is** detected, the device will not initialize. Proceed with the additional recovery steps listed under Recovery Mode below.

Recovery Mode

The device NVM image has exhibited damage or corruption.

Actions:

1. Wait 10 minutes for the NVM self-check process to complete. If during this period normal operation is achieved, the device will initialize and operate normally. No further action is required.
2. If after 10 minutes normal operation is *not* achieved:
    a. Check the System Event log for Windows OSs or driver message and kernel logs for Linux and ESXi based distributions. Recovery Mode is confirmed by presence of message/log entries as listed in the Firmware Recovery Mode Errors and Messages section above.
    b. Reboot the system and proceed with the additional recovery steps listed under NVM Image Restoration below.

Note:

- While in Recovery Mode, for Windows OSs, clicking on the adapter in Device Manager may present a dialog box indicating that Firmware Recovery Mode is active.
    - Once the dialog is dismissed, while the device appears to be functioning normally, it is in fact limited to only enable NVM image recovery.
- If the system is rebooted (versus power cycled), the driver status may not show a Code 10/yellow bang status as expected. Refer to events logged in System Event log for Windows OSs or driver message and kernel logs for Linux and ESXi based distributions to accurately assess the adapter status.
- When the adapter is in recovery mode, the link LED will not be lit and the adapter will not appear in the following locations:
    - F2 System Setup > Device Settings
    - System BIOS as a device for PXE Boot in UEFI boot mode

NVM Image Restoration

At this point, the device is in Firmware Recovery mode and its functionality is limited to only supporting restoration of the NVM image.

Actions:

1. 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. Please consult your operating system specific instructions on how to scan and repair potentially damaged system files.
2. 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.
3. After restoring the NVM image, perform an A/C power cycle of the system. Details for this are in the Other General Notes section below.

Note:

User configured settings (i.e. iSCSI target information, user-defined port/alternate MAC addresses) will not be restored to pre-recovery mode values.

Other General Notes

To perform an AC power cycle, do the following:

1. Shut down the system if it is powered up.
2. Unplug all AC power cords from the system.
3. Leave the AC power cords unplugged for 15 seconds to allow the system power supply to discharge completely.
4. Plug in AC power cords to the system.

Using Devlink to Update a Device's NVM

On a Linux system, when you update the NVM on some devices, the update may use the **devlink** interface, rather than the **ethtool** interface. This will happen if the following are true:

- You are updating an Intel Ethernet 800 Series device.
- Your system is running a distro that supports the devlink dev flash command.
- The firmware currently installed on the device supports it.
- The new NVM conforms to the correct PLDM format.

Most of the functionality and commands are the same with the following exceptions:

- You cannot update a device in Recovery Mode. (To update a device in recovery mode, you must download and install the Intel Ethernet driver set. See Recovery Mode for more information.)

- You cannot update the OROM or Netlist as a separate update, only as part of a full NVM update.
- If you specified a preservation level of PRESERVE_ALL, the system will immediately perform an EMPR reset after the NVM update.

On devices that support it, you can also use the **devlink** command line directly to update the device NVM:

```
devlink dev flash pci/0000:3b:00.0 file filename.bin
```

Where:

pci/0000:3b:00.0

The device you wish to update. You can get a list of devices with the devlink dev info command.

filename.bin

The file that contains the new NVM image.

However, we recommend that you use the NVMUpdate tool. See Intel® Ethernet NVM Update Tool for more information on this tool.

Firmware Logging

Intel® Ethernet 800 Series devices allow you to generate firmware logs for supported categories of events, to help debug complex issues with Customer Support. Firmware logging is enabled by default.

Note:

- Both the device and the driver need to support firmware logging for the functionality to work. If you are not able to set the configuration and the problem persists, reinstall the driver.
- You must have the latest base driver and NVM installed.
- Firmware logging events and levels are device-wide settings. Changing the log configuration on one port will apply it to all ports on a device.

Capturing a Firmware Log

To capture a firmware log, you must do the following:

1. Set the configuration for the firmware log. See later in this page for more information.
2. Perform the necessary steps to generate the issue you're trying to debug.
3. Capture the firmware log. (Exact steps will vary by operating system.)
4. Stop capturing the firmware log.

5. Reset your firmware log settings as needed.
6. Work with Customer Support to debug your issue.

> **Note:**
>
> Firmware logs are generated in a binary format and must be decoded by Customer Support. Information collected is related only to firmware and hardware for debug purposes.

Configuring Settings for Firmware Logs

Firmware logs capture information about different categories of events (called "modules"). A module corresponds to a general category of functionality, such as link topology detection or manageability.

The device's NVM sets default verbosity levels for each module. You can change the verbosity level per module; refer to OS-Specific Information below for more details. You can set only one log level per module, and each level includes the verbosity levels lower than it.

Available verbosity levels are:

- 0 = none
- 1 = error
- 2 = warning
- 3 = normal
- 4 = verbose

If you see errors or suspect the issue could fall into the below categories, setting the firmware logs to capture more verbosity for the corresponding module(s) in the right column will provide more information in the firmware log.

Modules for Events

Category of Event	Corresponding Module
Initialization	Control
NVM	• NVM • Authentication • VPD

Category of Event	Corresponding Module
I/O	• I2C • SDP • MDIO
Link Management	• Link Management • Link Control Technology • Link Topology Detection
Rx	• Parser • Switch • ACL • Post
Tx	• Scheduler • Tx Queue Management
AQ Interface	• Admin Queue • HDMA
Manageability	Manageability
Protocols	• LLDP • DCBx
Infrastructure	• Watchdog • Task Dispatcher • General • IOSF • PF Registration • Module Versions

Category of Event	Corresponding Module
XLR	XLR
QoS	DCB
Diagnostics	• SyncE • Health
TimeSync	Time Sync

Tips for Firmware Logs

- Firmware logs are for the hardest issues to debug. If you are experiencing issues, refer to the following sections for preliminary methods to diagnose problems:
    - Firmware
    - Troubleshooting
    - Health Status Messages
- We generally do not recommend to capture firmware logs at all times. If you suspect an issue, set the module(s) for the suspected event to a higher verbosity level, capture the firmware log, and then stop the log.
- Collecting firmware logs should not materially impact performance or CPU utilization.
- In general, set the logging level to Verbose only for the configuration group(s) or module(s) you need to debug. Setting too many modules to Verbose can overrun the buffer.
- You can try writing logs to a remote location or an external storage device, if your disk is full or your system does not have sufficient storage.

OS-Specific Information

Linux

At a high level, do the following to capture a firmware log in Linux:

1. Set log levels. For example:

```
echo normal > /sys/kernel/debug/ice/0000:18:00.0/fwlog/modules/all
```

2. Turn on firmware logging:

```
echo 1 > /sys/kernel/debug/ice/0000:18:00.0/fwlog/enable
```

3. Perform the necessary steps to generate the issue you are trying to debug.
4. Turn off firmware logging:

```
echo 0 > /sys/kernel/debug/ice/0000:18:00.0/fwlog/enable
```

5. Save data to a file:

```
cat /sys/kernel/debug/ice/0000:18:00.0/fwlog/data > fwlog.bin
```

6. Work with Customer Support to debug your issue.

Refer to the README in the driver tarball for more information on configuring firmware logs.

FreeBSD

Refer to the README in the driver tarball for more information on configuring firmware logs.

> **Note:** In FreeBSD, the driver can register/unregister to receive events per PF.

Windows

In Windows, you use PowerShell and Ethernet Cmdlets for Intel® Ethernet to configure firmware logging and capture firmware logs.

At a high level, do the following to capture a firmware log in Windows:

1. Set the configuration for the firmware log, using the Set-IntelEthernetLogConfig cmdlet in PowerShell.
2. Perform the necessary steps to generate the issue you're trying to debug.
3. Start capturing the firmware log, using the Start-IntelEthernetLog cmdlet.
4. Stop capturing the firmware log, using the Stop-IntelEthernetLog cmdlet.
5. Work with Customer Support to decode your firmware log file and debug the issue.

> **Note:**
>
> - Firmware logs will be captured in the file you designated with Start-IntelEthernetLog.
> - To disable firmware logging, use the Disable-IntelEthernetLogConfig cmdlet. To verify that firmware logging is disabled, run the Get-IntelEthernetLogConfig cmdlet; its results should say "Disabled."

ESXi

In ESXi, use esxcfg-module to set the configuration for firmware logs. Firmware logs are printed to kernel logs, with the tag FWLOG; use dmesg or read the file at /var/log/vmkernel.log.

At a high level, do the following to capture a firmware log in ESXi:

> **Note:** Refer to ESXi Example Commands below for all commands and parameters.

1. Set the configuration for the firmware log, using esxcfg-module. The ESXi driver uses the following module parameters for firmware logging:

   - FWLogEnable: Enables firmware logging functionality on the designated PF (0 = Disable, 1 = Enable). Use commas to separate the values for each PF; the first value is for PF0, second for PF1, and so on.
   - FWLogEvents: Designates the firmware events to log, using a bitmask. Binary math is required to set.
   - FWLogLevel: Sets the verbosity level for the firmware event’s log.

1. Redirect the kernel log or dmesg to a separate file for capturing the firmware log.
2. Reboot the system for changes to take effect.
3. After the system has rebooted, perform the necessary steps to generate the issue you’re trying to debug.
4. Work with Customer Support to decode your firmware log file and debug the issue.

   > **Note:** Firmware logs will be captured in the file you designated in step 2.

ESXi Example Commands

Use the following commands in ESXi for tasks related to firmware logging:

- To enable firmware logging and set the verbosity level for your desired events:

  ```
  esxcfg-module icen -s 'FWLogEnable=<values> FWLogEvents=<bitmask> FWLogLevel=<value>'
  ```

  For example, to enable firmware logging on PF0 and set all events to log warning messages, use:

  ```
  esxcfg-module icen -s 'FWLogEnable=1,0,0,0,0,0,0,0 FWLogEvents=255 FWLogLevel=2'
  ```

- To show the current configuration of the firmware log parameters:

  ```
  esxcfg-module -g <driver name>
  ```

  > **Note:**
  > If firmware logging is disabled, the FWLogEnable parameter should say “0” (disabled).

- To show a description of module parameters for firmware logging:

```
esxcfg-module -i <driver name>
```

Note: Look for the parameters that begin with FWLog.

- To redirect the firmware log to a file:

```
tail -f /var/log/vmkernel.log > filename.log
```

- To disable firmware logging:

```
esxcfg-module icen -s 'FWLogEnable=0 FWLogEvents=0 FWLogLevel=0'
```

Debug Dump

Intel® Ethernet 800 Series devices support debug dump, which allows you to obtain runtime register values from the firmware for “clusters” of events and then write the results to a single dump file, for debugging complicated issues in the field.

This debug dump contains a snapshot of the device and its existing hardware configuration, such as switch tables, transmit scheduler tables, and other information. Debug dump captures the current state of the specified cluster(s) and is a stateless snapshot of the whole device.

Note:

- The contents of the debug dump are not human-readable. You must work with Customer Support to decode the file.
- Debug dump is per device, not per PF.
- Debug dump writes all information to a single file.

Exact steps will vary by OS, but do the following to generate a debug dump log file:

1. Using the method appropriate for your OS (see OS-Specific Information below), specify the clusters for which you want to dump the hardware configuration. Supported clusters will vary by OS and hardware family.
2. Specify the path and filename for the dump file to be written to (optional depending on your OS).
3. Execute the command to write the debug dump file.
4. After the log file is written, work with Customer Support to decode the dump file.

OS-Specific Information

Use the following tools or commands to write the debug dump results to a dump file.

- Windows Server, Azure Stack HCI: Use either of the following:
  - `Write-IntelEthernetDebugDump` cmdlet in PowerShell (available with Ethernet Cmdlets for Intel® Ethernet; see the cmdlet help for more information)
  - Intel® Ethernet Inspector
- Windows: Not supported
- Linux: Use **debugfs** (see the Linux base driver README for more information)
- ESXi: Use **esxcli** (see below)
- FreeBSD: Use **sysctl** (see the FreeBSD base driver README for more information)

ESXi

> **Note:**
>
> For this functionality to work, you must have installed version 1.10.x or higher of the intnet tool, which is a plugin to the esxcli tool. You can download the latest version from the Intel Download Center here.

In esxcli, use the following command to generate the debug dump file for your specified cluster(s):

```
esxcli intnet debug fw dump <Cmd options>
```

Where `<Cmd options>` are:

-n, --vmnic <*string*>

Specifies the vmnic name to operate on. This field is required.

-c, --clusters <*string*>

[Optional] Specifies the clusters to dump.

To specify multiple clusters, enclose a single string in quotes, separated by commas with no spaces. For example:

```
esxcli intnet debug fw dump -n vmnic0 --clusters "ACL,L2P"
```

If `-c` is not specified, the driver dumps all clusters.

-l, --list

Displays the complete list of valid clusters on the screen.

To show the complete list of valid clusters, use the following:

```
esxcli intnet debug fw dump -n <vmnicX> -l
```

esxcli will output the debug dump results to a single file in the /scratch/core directory. The file naming convention is *vmnicX-<time-stamp>-dump.bin*, where vmnicX is the VMware device alias of the affected device.

Firmware Link Layer Discovery Protocol (FW-LLDP)

Devices based on the Intel® Ethernet 800 and 700 Series use a Link Layer Discovery Protocol (LLDP) agent that runs in the firmware. When it is running, it prevents the operating system and applications from receiving LLDP traffic from the network adapter.

- The FW-LLDP setting is per port and persists across reboots.
- The FW-LLDP Agent is required for Data Center Bridging (DCB) to function.

Adapters Based on the Intel® Ethernet 800 Series

FW-LLDP is disabled in NVM by default. To enable/disable the FW-LLDP Agent:

- Linux: Use ethtool to persistently set or show the fw-lldp-agent private flag.
- FreeBSD: Use sysctl to persistently set or show the fw_lldp_agent flag.
- ESX: Use the esxcli command to persistently set or get the fw-lldp-agent setting.
- Microsoft Windows: The base driver does not persistently change FW-LLDP. Use the LLDP Agent attribute in UEFI HII to persistently change the FW-LLDP setting. If you enable DCB when FW-LLDP is disabled, the base driver temporarily starts the LLDP Agent while DCB functionality is enabled.

Adapters Based on the Intel® Ethernet 700 Series

FW-LLDP is enabled in NVM by default. To enable/disable the FW-LLDP Agent:

- Linux: Use ethtool to set or show the disable-fw-lldp private flag.
- FreeBSD: Use sysctl to set or show the fw_lldp flag.
- ESX: Use the esxcfg-module command to set or get the LLDP module parameter.
- Microsoft Windows: Use the LLDP Agent attribute in UEFI HII to change the FW-LLDP setting. Note: You must enable the UEFI HII LLDP AGENT attribute for the FW-LLDP setting to take effect. If "LLDP AGENT" is set to disabled in UEFI HII, you cannot enable FW-LLDP from the OS.
- You must enable the LLDP Agent from UEFI HII to use DCB.

Flow Control

This feature 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. You may improve RDMA performance by enabling flow control on all nodes and on the switch to which they connect.

Note:

- For adapters to benefit from this feature, link partners must support flow control frames.
- On systems running a Microsoft Windows Server* operating system, enabling *QoS/priority* flow control will disable link level flow control.
- Some devices support Auto Negotiation. Selecting this will cause the device to advertise the value stored in its NVM (usually “Disabled”).

To change this setting in Intel® PROSet:

This setting is found on the PROset tab of the device’s Device Manager property sheet or in the Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

Possible values for this setting are:

- Disabled
- RX Enabled
- TX Enabled
- Auto-neegotiation (only available in some adapters)

Forward Error Correction (FEC) Mode

This feature 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.
- Intel® Ethernet 800 Series devices only enable Forward Error Correction (FEC) configurations that are supported by the connected media and which are expected to yield healthy Bit Error Rate (BER) connections.
  If you enable the registry keyword AllowNoFECModulesInAuto, Auto FEC negotiation will include *No FEC* in case your link partner does not have FEC enabled or is not FEC-

capable.

- To change this setting in Windows PowerShell, use the Set-IntelNetAdapterSetting cmdlet. For example:

  ```
  Set-IntelNetAdapterSetting -Name "<adapter_name>" -RegistryKeywordAllowNoFECModulesInAuto -RegistryValue 1
  ```

- To change this setting in Linux, use ethtool. For example:

  ```
  ethtool --set-priv-flags <ethX> allow-no-fec-modules-in-auto on
  ```

- If you are having link issues (including no link) at link speeds faster than 10Gbps, check your switch configuration and/or specifications. Many optical connections and direct attach cables require RS-FEC for connection speeds faster than 10Gbps. One of the following may resolve the issue:
  - Configure your switch to use RS-FEC mode.
  - Specify a 10Gbps, or slower, link speed connection.
  - If you are attempting to connect at 25Gbps, try using an SFP28 CA-S or CS-N Direct Attach cable. These cables do not require RS-FEC.
  - If your switch does not support RS-FEC mode, check with your switch vendor for the availability of a SW or FW upgrade.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Gigabit PHY Mode

This feature determines whether an adapter or link partner is designated as the primary. The other device is designated as the secondary. 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 primary. If both devices are multi-port devices, the one with higher seed bits becomes the primary.

The default setting is called *Hardware Default*. Setting the value to any other than the default will override the hardware default.

Note:

When Gigabit PHY Mode is forced to Primary mode on both the Intel adapter and its link partner, the link speed obtained by the Intel adapter may be lower than expected or link may not be established.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Gigabit PHY
 Mode"-DisplayValue "Auto Detect"
```

Possible values for this setting are:

- Force Primary Mode
- Force Secondary Mode
- Auto-Detect

Note:

When Gigabit PHY Mode is forced to Primary mode on both the Intel® Ethernet adapter and its link partner, the link speed obtained by the Intel adapter may be lower than expected or link may not be established.

Intel® Ethernet Flow Director

The Intel® Ethernet Flow Director (Intel® Ethernet FD) performs the following tasks:

- Directs receive packets according to their flows to different queues
- Enables tight control on routing a flow in the platform
- Matches flows and CPU cores for flow affinity
- Depending on the device family: Supports multiple parameters for flexible flow classification and load balancing (in SFP mode only)

Depending on the driver and device family, the driver might support the following flow types:

- IPv4
- TCPv4
- UDPv4
- SCTPv4

- IPv6
- TCPv6
- UDPv6
- SCTPv6

Each flow type supports valid combinations of IP addresses (source or destination) and UDP/TCP ports (source and destination). You can supply only a source IP address, a source IP address and a destination port, or any combination of one or more of these four parameters.

The following table summarizes supported Intel Ethernet Flow Director features across Intel Ethernet controllers.

Feature	500 Series	700 Series	800 Series
VF Flow Director	Supported	Routing to VF not supported	Not supported
IP Address Range Filter	Supported	Not supported	Field masking
IPv6 Support	Supported	Supported	Supported
Configurable Input Set	Configured per port	Configured globally	Configured per port
ATR	Supported	Supported	Not supported
Flex Byte Filter	Starts at beginning of packet	Starts at beginning of payload	Starts at beginning of packet
Tunneled Packets	Filter matches outer header	Filter matches inner header	Filter matches inner header

See the Linux driver READMEs for more information on configuring this feature.

Interrupt Moderation Rate

This feature 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Possible values for this setting are:

- Adaptive
- Extreme
- High
- Medium
- Low
- Minimal
- Off

Jumbo Frames

The Jumbo Frames feature enables or disables Jumbo Packet capability. The standard Ethernet frame size is 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® Ethernet 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 that was 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.
- Jumbo Frames are not supported over Intel® Advanced Network Services (Intel® ANS) VLANs under Microsoft Windows 10.

Note:

- 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in the Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Jumbo Pack
et"-DisplayValue "4088 Bytes"
```

Possible values for this setting are:

- Disabled (1514 bytes)
- 4088 Bytes
- 9014 Bytes

Note: Set the switch 4 bytes higher for CRC, plus 4 bytes if using VLANs.

Link State on Interface Down

This feature sets if link is enabled or disabled when the interface is brought down. If this is set to **Disabled** and you bring an interface down (using an administrative tool, or in another way), then the port will lose its link. This allows an attached switch to detect that the interface is no longer up. However, if Wake on LAN or manageability is enabled on this port, the link will remain up.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Link State onInterface Down" -DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

Locally Administered Address

The 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "LocallyAdministered Address" -DisplayValue "<desired address>"
```

Possible values for this setting are:

- 0000 0000 0001 - FFFF FFFF FFFD

Note: 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 FFFFUnicast (legal): 0070 4567 8999
```

Note: In a team, Intel PROSet uses either of the following:

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

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

Log Link State Event

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

The following events are logged:

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

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in the Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the `Set-IntelNetAdapterSetting` cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Log LinkSta
te Event" -DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

Low Latency Interrupts

Low Latency Interrupts (LLI) enable the network device to bypass the configured interrupt moderation scheme based on the type of data being received. LLI 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 that contain 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Low Latenc
yInterrupts" -DisplayValue "Port-Based"
```

Possible values for this setting are:

- Disable
- PSH Flag-Based
- Port-Based

Malicious Driver Detection (MDD) for VFs

Some Intel® Ethernet devices use Malicious Driver Detection (MDD) to detect malicious traffic from the VF and disable Tx/Rx queues or drop the offending packet until a VF driver reset occurs. You can view MDD messages in the PF's event log.

- If the device supports automatic VF resets and the driver detects an MDD event on the receive path, the PF will automatically reset the VF and reenable queues. If automatic VF resets are disabled, the PF will not automatically reset the VF when it detects MDD events. See the table below for supported MDD features.
- If the PF driver logs MDD events from the VF, confirm that the correct VF driver is installed.
- To restore functionality, you can manually reload the VF or VM or, if supported by the device, enable automatic VF resets.

The following table shows MDD capabilities by device family:

- Intel Ethernet 800 Series
- Intel Ethernet 700 Series
- Intel Ethernet 500 Series
- Intel I350 Gigabit Network Connection

Feature	**800 Series**	**700 Series**	**500 Series**	**I350**
Automatically resets the VF and reenables queues after MDD events	If enabled	If enabled	Yes	Yes
Can disable automatic VF reset after MDD events	Yes	Yes	No	No

MDD Auto Reset VFs

This feature automatically resets the virtual machine immediately after the device detects a Malicious Driver Detection (MDD) event on the receive path.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "MDD Auto Reset VFs"-DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

Max Number of RSS Queues Per Vport

This setting sets the maximum number of Receive Side Scaling (RSS) queue pairs per VF.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Max Numbe
r ofRSS Queues Per Vport" -DisplayValue "4 Queues"
```

Possible values for this setting are:

- 2 Queues
- 4 Queues
- 8 Queues
- 16 Queues

Offloads

- IPv4 Checksum Offload
- Large Send Offload (IPv4 and IPv6)
- NVGRE Encapsulated Task Offload
- QoS Offload
- TCP Checksum Offload (IPv4 and IPv6)
- UDP Checksum Offload (IPv4 and IPv6)
- UDP Segmentation Offload (IPv4 and IPv6)
- VXLAN Encapsulated Task Offload

In addition to the offloads listed above, see the following pages for related information:

- Priority and VLAN Tagging
- Virtual Machine Queue Offloading

IPv4 Checksum Offload

This setting 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 (on RX) and computation (on TX) for the operating system.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Possible values for this setting are:

- Disabled
- RX Enabled
- TX Enabled
- RX & TX Enabled

Large Send Offload (IPv4 and IPv6)

This setting 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Possible values for this setting are:

- Enabled
- Disabled

NVGRE Encapsulated Task Offload

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 NVGRE processing, offloading it from the operating system. This reduces CPU utilization.

To change this setting in Intel® PROSet:

This setting is found in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU). On the device’s Device Manager property sheet, it is found on the Advanced tab, under the Offloading Options > Encapsulated Task Offload setting.

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "NVGREEncapsulated Task Offload" -DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

QoS Offload

This setting configures the Quality of Service (QoS) offload for the miniport adapter. This feature allows you to set a bandwidth cap and reservation to one or more virtual machines on a physical device, including both software VMs and SR-IOV interfaces.

To change this setting in Intel® PROSet:

This setting is found on the Adapter tab and in the Adapter Settings panel of Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "QoS Offload "-DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

TCP Checksum Offload (IPv4 and IPv6)

This setting 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "TCP ChecksumOffload (IPv4)" -DisplayValue "Tx Enabled"
```

Possible values for this setting are:

- Disabled
- RX Enabled
- TX Enabled
- RX & TX Enabled

UDP Checksum Offload (IPv4 and IPv6)

This setting 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Possible values for this setting are:

- Disabled
- RX Enabled
- TX Enabled
- RX & TX Enabled

UDP Segmentation Offload (IPv4 and IPv6)

This setting allows the adapter to segment UDP packets with payloads up to 64K into valid Ethernet frames. Because 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 may use fewer CPU resources.

With Offloading off, the operating system segments UDP packets into valid Ethernet frames.

With Offloading on, the adapter segments UDP packets for the operating system.

Note: UDP Segmentation Offload requires:

- Microsoft* Windows Server* 2019, Version 1903, or later
- Linux* kernel 4.18, or later

To change this setting in Intel® PROSet:

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "UDP SegmentationOffload (IPv4)" -DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

VXLAN Encapsulated Task Offload

Virtual Extensible LAN (VXLAN) allows you to extend an L2 network over an L3 network, which may be useful in a virtualized or cloud environment. Some Intel® Ethernet devices perform

VXLAN processing, offloading it from the operating system. This reduces CPU utilization.

VXLAN may be useful in multi-tenant environments such as cloud service providers where the number of VLANs exceeds the 4094 limit imposed by the 12-bit VLAN ID used in Ethernet data frames.

To change this setting in Intel® PROSet:

This setting is found in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU). On the device's Device Manager property sheet, it is found on the Advanced tab, under the Offloading Options > Encapsulated Task Offload setting.

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "VXLAN En
capsulatedTask Offload" -DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

Performance Options

Optimizing Performance

You can configure advanced settings on Intel network adapters to help optimize server performance. This section provides tips for:

- General Optimization
- Optimization for Specific Usage Models

Note:

- Linux users: See the README file in the Linux driver package for Linux-specific performance enhancement details.

- 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.
- 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.

- For an Intel Ethernet 700 Series 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.
- 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.

Jumbo Frames

Enabling jumbo frames may increase throughput. You must enable jumbo frames on all of your network components to get any benefit.

RSS Queues

If you have multiple 10Gbps (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 Non-Uniform Memory Access (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.For Microsoft Windows systems, do the following:

1. Identify the adapter ports to be adjusted and inspect 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 BaseProcessor Number" -DisplayValue <RSS Base Proc Value>
   ```

4. Use the Get-NetAdapterAdvancedproperty cmdlet to check that the right values have been set:

   ```
   Get-NetAdapterAdvancedproperty -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.

CPU Affinity

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.

Optimization for Specific Usage Models

The following sections describe possible tasks you can try to optimize performance for specific server usage models.

Optimize for Quick Response and Low Latency

Useful for: Video, audio, and High Performance Computing Cluster (HPCC) servers

- Minimize or disable interrupt moderation rate.
- Disable offload TCP segmentation.
- Disable jumbo packets.
- Increase transmit descriptors.
- Increase receive descriptors.
- Increase RSS queues.

Optimize for Throughput

Useful for: Data backup/retrieval and file servers

- 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.

Optimize for CPU Utilization

Useful for: Application, web, mail, and database servers

- 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.

Tuning Performance with SR-IOV

When SR-IOV is enabled in Hyper-V, the following steps can help to improve performance between VM to VM and VM to Host.

From host OSes on both Host 1 and Host 2:

1. Enable RSS on the PF and vSwitch:

   ```
   Enable-NetAdapterRss -name "ADAPTER_NAME"
   ```

2. Enable 4 queues per VF:

   ```
   Set-VMNetworkAdapter -VMName "YOUR_TEST_VM_NAME" -IovQueuePairsRequested 4Get-VmNetworkAdapter -VMName * | where {$_.SwitchName -eq "YOUR_TEST_SWITCH_NAME"} |  Set-VmNetworkAdapter -IovQueuePairsRequested 4
   ```

3. Ensure the VMs have at least twice as many vCPUs as RSS queues. In this case, set the number of total processors in the VM to 8. To do this:
   a. Turn off the VM.
   b. In the VM, click **Settings**.
   c. Under **Hardware**, select **Processor**.
   d. Change the value of **Number of virtual processors** to 8.
   e. Apply the change.
4. For Windows Server 2022, issue the following command while the VM is in the off state:

```
Set-VMProcessor -VMName "YOUR_VM_Name" -HwThreadCountPerCor
e 1 -Count 8
```

In both guest OSes VM1 and VM2:

1. Set RSS queues to 4 for all VFs in the guest OSes:

```
Set-NetAdapterRss -InterfaceDescription *adaptive* -NumberOfReceiveQue
ues 4
```

2. Update the number of queues in the guest OS:

```
Set-NetAdapterAdvancedProperty -Name "your_adapter_name_from_guest_
os" -DisplayName"Maximum Number of RSS Queues" -DisplayValue "8 Qu
eues"
```

> **Note:**
>
> In the locations where there are settings for the number of queues, that value can be anything from 1 to 16. If you want more total throughput, increase the number of queues. When updating the number of queues, you **must** set IovQueuePairsRequested to a value that is equal to or greater than the number of queues you want to use in the VM.

Transmit Balancing

Some Intel® Ethernet 800 Series devices allow you to enable a transmit balancing feature to improve transmit performance under certain conditions. When the feature is enabled, you should experience more consistent transmit performance across queues and/or PFs and VFs.

By default, transmit balancing is disabled in the NVM. To enable this feature, use one of the following to persistently change the setting for the device:

- Use the Ethernet Port Configuration Tool (EPCT) to enable the tx_balancing option. Refer to Ethernet Port Configuration Tool (EPCT) for more information.
- Enable the Transmit Balancing device setting in UEFI HII.
- Enable transmit balancing via Linux devlink. Refer to the Linux readme inside the driver tarball for more information.

When the driver loads, it reads the transmit balancing setting from the NVM and configures the device accordingly.

> **Note:**
>
> - The user selection for transmit balancing in EPCT, HII, or Linux devlink is persistent across reboots. You must reboot the system for the selected setting to take effect.

- This setting is device wide.
- The driver, NVM, and DDP package must all support this functionality to enable the feature.

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.

To change this setting in Intel® PROSet:

You must install Intel PROSet to use Performance Profiles.

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Profile" -DisplayValue "Standard Server"
```

Possible values for this setting are:

- 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.

Note:

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

Teaming Considerations

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

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

Note:

This feature is not configurable through Intel PROSet ACU. On Microsoft Windows Server 2019, Microsoft Windows* 10 Version 1809, and later, use Windows PowerShell.

See Adapter Teaming for more information on teams.

Power Options

The Power Management tab in the device's Device Manager property sheet or the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU) includes several settings that control the device's power consumption. For example, you can set the adapter to reduce its power consumption if the cable is disconnected.

See the following for more information on power options:

- Wake on LAN (WoL) Options
  - WoL Supported Devices
- Other Power Options
  - Wake from S0ix on Magic Packet
  - Reduce Power if Cable Disconnected & Reduce Link Speed During Standby
  - Ultra Low Power Mode When Cable is Disconnected
  - Selective Suspend
  - Selective Suspend Idle Timeout
  - Energy Efficient Ethernet

ACPI Power States

Advanced Configuration and Power Interface (ACPI) supports a variety of power states. Each state represents a different level of power, from fully powered up to completely powered down, with partial levels of power in each intermediate state.

The ACPI power states are:

S0:

On and fully operational

S1:

System is in low-power mode (sleep mode). The CPU clock is stopped, but RAM is powered on and being refreshed.

S2:

Similar to S1, but power is removed from the CPU.

S3:

Suspend to RAM (standby mode). Most components are shut down. RAM remains operational.

S4:

Suspend to disk (hibernate mode). The memory contents are swapped to the disk drive and then reloaded into RAM when the system is awakened.

S5:

Power off

Microsoft Windows Server* is ACPI-capable. It 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 Ethernet 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.

Wake on LAN (WoL) Options

The ability to remotely wake computers is an important development in computer management. This feature has evolved from a simple remote power-on capability to a complex system interacting with a variety of device and operating system power states.

The Intel® PROSet Power Management tab or the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU) 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 or the Adapter Settings panel in Intel PROSet ACU includes a setting called Wake on Magic Packet from power off state. Enable this setting to explicitly allow wake-up with a Magic Packet from shutdown under APM power management mode.

Note:

- To use the Wake on Directed Packet feature, WoL must first be enabled in the EEPROM using BootUtil. See Intel® Ethernet Flash Firmware Utility for more information on 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

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

Intel® Ethernet 800 Series

The following devices do not support WoL:

- Intel® Ethernet Network Adapter E810-2C-Q2
- Intel® Ethernet Network Adapter E810-C-Q2
- Intel® Ethernet Network Adapter E810-C-Q2T
- Intel® Ethernet Network Adapter E810-XXV-4
- Intel® Ethernet Network Adapter E810-XXV-4T
- Intel® Ethernet Network Adapter E810-XXV-2

Intel® Ethernet 700 Series

The following devices do not support WoL:

- Intel® Ethernet Network Adapter I710-T4L for OCP 3.0
- Intel® Ethernet Network Adapter I710-T4L
- 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
- Intel® Ethernet Network Adapter X710-T2L
- Intel® Ethernet Network Adapter X710-T4L
- Intel® Ethernet Network Adapter X710-TL

Intel® Ethernet 500 Series

The following devices do not support WoL:

- 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

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

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

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

Intel® Ethernet 300 Series

The following 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

Other Power Options

Wake from S0ix on Magic Packet

Enables this device to bring the system out of an S0ix power state when the device receives a Magic Packet.

Possible values for this setting are:

- Enabled
- Disabled

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.

Possible values for this setting vary 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.

Selective Suspend

Enables the device to enter a low power state when the network is idle.

To change this setting in Intel PROSet:

Possible values for this setting are:

- Enabled
- Disabled

Selective Suspend Idle Timeout

Sets the length of time that the network is idle before the device enters a low power state. You must enable Selective Suspend for this setting to take effect.

To change this setting in Intel PROSet:

Possible values for this setting are:

- 1 - 60 in seconds

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.

Note:

- Both ends of the EEE link must automatically negotiate link speed.
- EEE is not supported on every adapter.

Priority and VLAN Tagging

This setting enables the adapter to offload the insertion and removal of priority and VLAN tags for transmit and receive.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "DCB"-Displ
ayValue "Disabled"Set-IntelNetAdapterSetting -Name "<adapter_name>" -Display
Name "PacketPriority & VLAN" -DisplayValue "VLAN Enabled"
```

Possible values for this setting are:

- Priority & VLAN Disabled
- Priority Enabled
- VLAN Enabled
- Priority & 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.

To implement QoS, the adapter must be connected to a switch that supports and is configured for 802.1p QoS. Priority-tagged frames allow programs that deal with real-time events to make the most efficient use of network bandwidth. High priority packets are processed before lower priority packets.

To change this setting in Intel® PROSet:

Tagging is enabled and disabled using the following fields:

- **In Windows Server*:** Use the “QoS Packet Tagging” field in the Advanced tab in Intel PROSet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).
- **In other versions of Windows*:** Use the “Priority/VLAN Tagging” setting on the Advanced tab in Intel PROSet or in the Adapter Settings panel in Intel PROSet ACU.

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"-Displ
ayValue "Disabled"Set-IntelNetAdapterSetting -Name "<adapter_name>" -Display
Name "PacketPriority & 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 Windows Server, have a utility for 802.1p packet prioritization. For more information, see the Windows system help and Microsoft's knowledge base.

Receive Buffers

This setting 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Receive Buf
fers"-DisplayValue "256"
```

Possible values for this setting are:

- 128-4096, in intervals of 64, for all adapters

We recommend the following values:

- Teamed adapter: 256
- Using IPSec and/or multiple features: 352

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.

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

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 preferred 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.
- **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, or faster, adapters.
- **Preferred NUMA Node.** This setting allows you to choose the preferred NUMA (Non-Uniform Memory Access) node to be used for memory allocations made by the network adapter. In addition, the system will attempt to use the CPUs from the preferred NUMA node first for the purposes of RSS. On NUMA platforms, memory access latency is dependent on the memory location. Allocation of memory from the closest node helps improve performance. The Windows* Task Manager shows the NUMA Node ID for each processor.

  Note:

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

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Possible values for this setting are:

- 1 queue is used when low CPU utilization is required.
- 2 queues are used when good throughput and low CPU utilization are required.
- 4 or more queues are used for applications that demand maximum throughput and transactions per second.

Note:

- Not all settings are available on all adapters.
- 8, or more, queues are only available when Intel® PROSet for Windows* Device Manager or Intel PROSet ACU is installed. If Intel PROSet is not installed, only 4

queues are available.

- Using 8 or more queues requires the system to reboot.

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.

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.

See the following subsections for more information.

- Flash Images
- Intel® Boot Agent

Note:

See Configuring the UEFI Network Stack for PXE for information on enabling a PXE network boot.

Enabling 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. Do the following:
  1. From the command line, enter the following to enumerate the ports available in your system:

     ```
     BOOTUTIL -E
     ```

  2. Choose a port.
  3. Enter the following to select the port you wish to enable. For example:

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

See Intel® Ethernet Flash Firmware Utility for more information on using BootUtil.

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 Ethernet adapter or on the system board.

This page describes how to update the flash for the following OSes:

- Microsoft Windows*
- Linux*
- UEFI

Note:

You cannot update the flash of a device in the “Pending Reboot” state. Reboot your system before attempting to update the device’s flash.

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. See Intel® Ethernet Flash Firmware Utility for more information on using 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.

Note:

This functionality is not supported in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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.

Note:

This functionality is not supported in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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 at the command line:

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

See Intel® Ethernet Flash Firmware Utility for more information on using BootUtil.

> **Note:**
>
> 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.

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
```

See Intel® Ethernet Flash Firmware Utility for more information on using BootUtil.

> **Note:**
>
> 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.

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 the BOOTP protocol.

Configuring the Intel Boot Agent Client

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 preboot environment (before operating system is loaded)

The Intel Boot Agent supports PXE in preboot and Microsoft Windows environments. In each of these environments, a single user interface allows you to configure PXE protocols on Intel® Ethernet adapters.

Important:

If spanning tree protocol is enabled on a switch port through which a port is trying to use PXE, the delay before the port starts forwarding can cause a DHCP timeout. Either disable spanning tree or turn on the feature that allows the port to begin forwarding of packets immediately ("port fast learning" for Cisco switches), rather than wait until the spanning tree discovery is complete.

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 a Preboot 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 preboot (operating system independent) configuration setup program contained within the adapter's flash ROM. You can access this preboot 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:

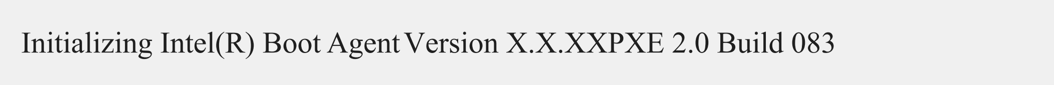

```
Initializing Intel(R) Boot Agent Version X.X.XXPXE 2.0 Build 083
```

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 space bar 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 following describes the available configuration settings and their possible values.

Network Boot Protocol:

Select PXE for use with Network management programs, such as LANDesk* Management Suite. Depending on the configuration of the Intel Boot Agent, this parameter may not be changeable.

Possible values include:

- PXE (Preboot eXecution Environment)

Boot Order:

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**.

Possible values include:

- Use BIOS Setup Boot Order
- Try network first, then local drives
- Try local drives first, then network
- Try network only
- Try local drives only

Legacy OS Wakeup Support:

(For 82559-based adapters only)

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.

Possible values include:

- 0 = Disabled (Default Value)
- 1 = Enabled

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 Preboot PXE Environments below.

Setting Up the Intel Boot Agent Target/Server

For the Intel Boot Agent software to perform its intended job, a server must be 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 in the software release. 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 Troubleshooting

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 Intel Boot Agent in a Preboot 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 Customer Support if you need further assistance.

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 Preboot PXE Environments

Anytime the configuration setup menu is displayed (see Configuring the Intel Boot Agent in a Preboot PXE Environment above), 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.

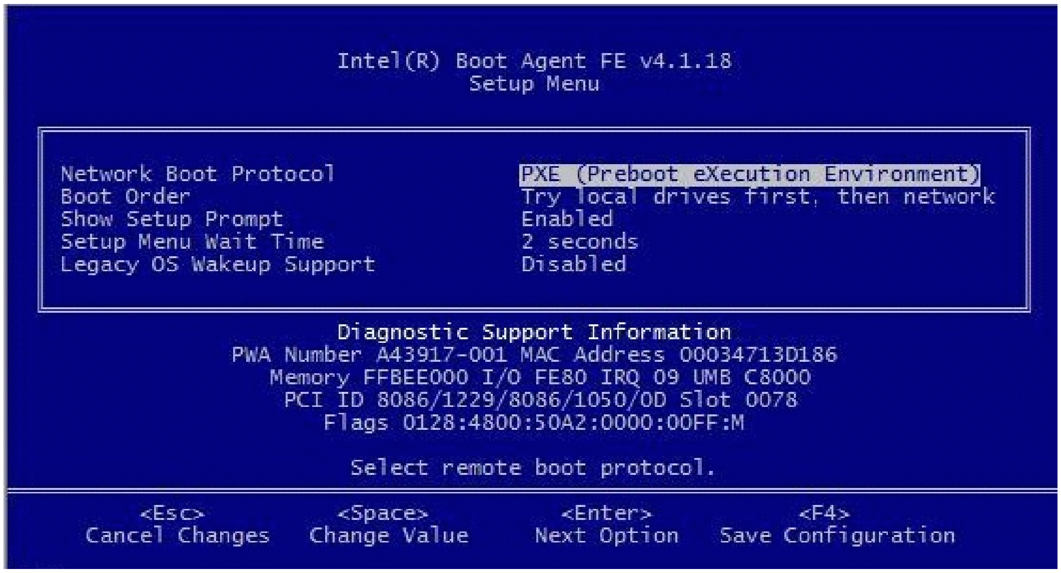

```
                    Intel(R) Boot Agent FE v4.1.18
                              Setup Menu

 Network Boot Protocol              PXE (Preboot eXecution Environment)
 Boot Order                         Try local drives first, then network
 Show Setup Prompt                  Enabled
 Setup Menu Wait Time               2 seconds
 Legacy OS Wakeup Support           Disabled

                    Diagnostic Support Information
             PWA Number A43917-001 MAC Address 00034713D186
               Memory FFBEE000 I/O FE80 IRQ 09 UMB C8000
                PCI ID 8086/1229/8086/1050/0D Slot 0078
                  Flags 0128:4800:50A2:0000:00FF:M

                     Select remote boot protocol.

      <Esc>            <Space>            <Enter>            <F4>
  Cancel Changes    Change Value       Next Option     Save Configuration
```

Diagnostics information may include the following items:

- 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.

Remote Direct Memory Access (RDMA)

Remote Direct Memory Access, or RDMA, allows a network device to transfer data directly to and from application memory on another system, increasing throughput and lowering latency in certain networking environments.

- Intel® Ethernet 800 Series devices support both iWARP and RoCEv2.
- Intel® Ethernet X722 Series devices only support iWARP.

The major difference is that iWARP performs RDMA over TCP, while RoCEv2 uses UDP.

On devices with RDMA capabilities, RDMA is supported on the following operating systems:

- Linux*
- FreeBSD*
- VMware* ESXi*
- Microsoft* Windows Server*

To avoid performance degradation from dropped packets, enable link level flow control or priority flow control on all network interfaces and switches.

Note:

- On systems running a Microsoft Windows Server operating system, enabling *QoS/priority flow control will disable link level flow control.
- Devices based on the Intel Ethernet 800 Series do not support RDMA when operating in multiport mode with more than 4 ports.
- On Linux systems, RDMA and link aggregation (LAG, also known as *bonding*) are not compatible on most devices. If RDMA is enabled, bonding will not be functional.
    - On Intel Ethernet 810 Series devices, RDMA and LAG are compatible if all the following are true:
        - □ You are using an Intel Ethernet 810 Series device with the latest drivers and NVM installed.
        - □ RDMA technology is set to RoCEv2.
        - □ LAG configuration is either active-backup or active-active.
        - □ Bonding is between two ports within the same device.
        - □ The QoS configuration of the two ports matches prior to the bonding of the devices.

      See the README inside the Linux ice driver tarball for more information.

RDMA on Linux or FreeBSD

For Intel Ethernet devices that support RDMA on Linux or FreeBSD, use the drivers shown in the following table.

Device	**Linux**		**FreeBSD**		**S upported P rotocols**
	Base Driver	RDMA Driver	Base Driver	RDMA Driver	
Intel Ethernet 800 Series	ice	irdma	ice	irdma	RoCEv2, iWARP
Intel Ethernet X722 Series	i40e	irdma	ixl	not s upported	iWARP

Basic Installation Instructions

At a high level, installing and configuring RDMA on Linux or FreeBSD consists of the following steps. See the README file inside the appropriate RDMA driver tarball for full details.

1. Install the base driver.
2. Install the RDMA driver.
3. Install and patch any user-mode RDMA libraries. Exact steps will vary by operating system; refer to the RDMA driver readme for details.
4. Enable flow control on your device. Refer to the base driver README for details and supported modes.
5. If you are using RoCE, enable flow control (PFC or LFC) on the device and endpoint your system is connected to. See your switch documentation and, for Linux, the Intel® Ethernet 800 Series Linux Flow Control Configuration Guide for RDMA Use Cases for details.

RDMA for Virtualized Environments in Linux

Devices based on the Intel Ethernet 800 Series support RDMA in a Linux VF on supported Windows or Linux hosts. Refer to the README file inside the Linux RDMA driver tarball for more information on how to load and configure RDMA in a Linux VF.

RDMA on Microsoft Windows

RDMA for Network Direct (ND) User-Mode Applications

Network Direct (ND) allows user-mode applications to use RDMA features.

> **Note:**
>
> User-mode applications may have prerequisites such as Microsoft HPC Pack or Intel MPI Library, refer to your application documentation for more details.

Supported Operating Systems

Intel® Ethernet User Mode RDMA Provider is supported on Microsoft Windows Server.

RDMA User Mode Installation

Follow the steps below to install user-mode Network Direct features:

1. Download the software package you want for the release. See Install Windows* Drivers for more information.
   a. If you are installing via the complete driver pack:
      i. In the extracted files from the download, navigate to \APPS\PROSETDX and then the Windows subfolder corresponding to your version of Windows (32-bit or 64-bit).
      ii. Inside the Winx64 or Win32 folder, double-click on *DxSetup.exe* to launch the install wizard.
   b. If you are installing via the separate webpacks for base drivers and Intel® PROSet:
      i. Download and extract the webpack for Intel PROSet.
      ii. In the extracted files, double-click on the **.exe** file to launch the install wizard.
2. On the Setup Options screen, select **Intel® Ethernet User Mode RDMA Provider**.
3. On the RDMA Configuration Options screen, select **Enable RDMA routing across IP Subnets** if desired. Note that this option is displayed during base driver installation even if user mode RDMA 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 RDMA Port Mapping Service rule in Windows Firewall** and the networks to which to apply the rule.

   > **Note:**
   > If Windows Firewall is disabled or you are using a third party firewall, you will need to add this rule manually.

5. Continue with driver and software installation.

RDMA Network Direct Kernel (NDK)

RDMA Network Direct Kernel (NDK) functionality is included in the Intel base networking drivers and requires no additional features to be installed.

RDMA Routing across IP Subnets

If you want to allow NDK's RDMA functionality across subnets, you will need to select **Enable RDMA routing across IP Subnets** on the RDMA Configuration Options screen during base driver installation. See Install Windows* Drivers for basic installation instructions.

Enabling Priority Flow Control (PFC) on a Microsoft Windows Server

To avoid performance degradation from dropped packets, enable priority flow control (PFC) or link level flow control on all network interfaces and switches.

Note:

On systems running a Microsoft Windows Server operating system, enabling *QoS/priority flow control will disable link level flow control.

Use the following PowerShell* commands to enable PFC on Microsoft Windows Server:

```
Install-WindowsFeature -Name Data-Center-Bridging -IncludeManagementToolsNew-NetQoSPolicy "SMB" -NetDirectPortMatchCondition 445 -PriorityValue8021Action 3Enable-NetQosFlowControl -Priority 3Disable-NetQosFlowControl -Priority 0,1,2,4,5,6,7New-NetQosTrafficClass -Name "SMB" -Priority 3 -BandwidthPercentage 60 -Algorithm ETSSet-NetQosDcbxSetting -Willing $FALSEEnable-NetAdapterQos -Name "Slot1 4 2 Port 1"
```

Verifying RDMA Operation with Microsoft PowerShell

You can check that RDMA is enabled on the network interfaces using the following Microsoft PowerShell command:

```
Get-NetAdapterRDMA
```

Use the following PowerShell command to check if the network interfaces are RDMA capable and multichannel is enabled:

```
Get-SmbClientNetworkInterface
```

Use the following PowerShell command to check if Network Direct is enabled in the operating system:

```
Get-NetOffloadGlobalSetting | Select NetworkDirect
```

Use **netstat** to make sure each RDMA-capable network interface has a listener at port 445 (Windows Client OSs that support RDMA may not post listeners). For example:

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

RDMA for Virtualized Environments in Windows

To enable RDMA functionality on virtual adapter(s) connected to a VMSwitch, you must:

- Enable SR-IOV (Single Root IO Virtualization) and VMQ (Virtual Machine Queues) advanced properties on each port.
- Set the number of VFs to enable with RDMA capabilities. You can enable up to 32 VFs with RDMA capabilities.

Under certain circumstances, you may disable these settings by default. You can manually set these options in the Adapter Settings panel of Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU), in the Advanced tab of the adapter properties dialog box, or with the following PowerShell commands:

```
Set-NetAdapterAdvancedProperty -Name <nic_name> -RegistryKeyword *SRIOV
-RegistryValue 1Set-NetAdapterAdvancedProperty -Name <nic_name> -RegistryK
eyword *VMQ -RegistryValue 1Set-NetAdapterAdvancedProperty -Name <nic_na
me> -RegistryKeyword RdmaMaxVfsEnabled -RegistryValue <1-32>
```

Configuring RDMA Guest Support (NDK Mode 3)

NDK Mode 3 allows kernel mode Windows components to use RDMA 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 setting on the device.
3. Enable SR-IOV on the VMSwitch bound to the device by performing the following for all physical 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 setting. 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 physical functions (PFs) and other VFs:

   ```
   Set-NetAdapterAdvancedProperty -Name <device_name> -RegistryKeywor
   d 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-NetAdapterGet-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 Intel 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 1Get-NetAdapterRdma | Enable-NetAdapterRdma
    ```

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

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

Testing NDK: Microsoft Windows SMB Direct with DiskSPD

This section outlines the recommended way to test 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 and connect at least two servers running a supported Microsoft Windows Server operating system, with at least one RDMA-capable Intel Ethernet device 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. You can use the following PowerShell command:

   ```
   New-SmbShare -Name <SMBsharename> -Path <SMBsharefilepath> -FullA
   ccess  <domainname>\Administrator,Everyone
   ```

   For example:

   ```
   New-SmbShare -Name RAMDISKShare -Path R:\RAMDISK -FullAccess gr
   oup\Administrator,Everyone
   ```

3. Download and install the Diskspd Microsoft utility from here: https://gallery.technet.microsoft.com/DiskSpd-a-robust-storage-6cd2f223
4. Using CMD or Powershell, use the **cd** command to change to the DiskSpd folder, and then run tests. (Refer to Diskspd documentation for more details on parameters)

   For 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 \\<SMBserverTestIP>\<SMBsharename>\test.dat:

   ```
   .\diskspd.exe -b4K -d60 -h -L -o16 -t16 -r -w0 -c10G \\<SMBserverTestIP>\
   <SMBsharename>\test.dat
   ```

5. Verify that RDMA traffic is running using perfmon counters such as “RDMA Activity” and “SMB Direct Connection”. Refer to Microsoft documentation for more details.

RDMA Windows Performance Monitoring

You can use **perfmon**, or other performance monitoring tool, to monitor and display RDMA counters and statistics. Refer to Microsoft documentation for more details. Use the Register-IntelEthernetRDMACounterSet cmdlet to register the RDMA statistics counters for the specific device with perfmon. Refer to Configuring Features with Windows PowerShell* for more information about how to install and use Intel cmdlets. You can use the following PowerShell command to register the RDMA statistics for all supported devices:

```
Register-IntelEthernetRDMACounterSet
```

You can use the following PowerShell cmdlet to unregister the RDMA statistics:

```
Unregister-IntelEthernetRDMACounterSet
```

Accessing Remote NVM Express* Drives Using RDMA

RDMA provides a high throughput, low latency means to directly access NVM Express* (NVMe*) drives on a remote server.

Refer to the following for details on supported operating systems and how to set up and configure your server and client systems:

- NVM Express over TCP for Intel® Ethernet Products Configuration Guide
- NVM Express over Fabrics for Intel® Ethernet Products with RDMA Configuration Guide

Both guides are available on the Intel Technical Library.

Setting Speed and Duplex

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. 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.

Note:

- Only experienced network administrators should force speed and duplex manually.
- Fiber-based adapters operate only in full duplex at their native speed. You cannot change the speed or duplex of Intel adapters that use fiber cabling.
- Some devices may list 10Mbps and 100Mbps in full or half duplex as options. These settings are not recommended.
- Link speed information in Intel® PROSet 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.
- Auto-negotiation and Auto-Try are not supported on devices based on the Intel® Ethernet Connection X552 and Intel® Ethernet Connection X553.

Manually Configuring Duplex and Speed Settings

Note:

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.

In Intel® PROSet for Windows* Device Manager:

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**.

In Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU), link speed is reported on the Adapter Information panel. Change speed and duplex in the Adapter Settings panel.

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"
```

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 or in the Adapter Information panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU). 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.

Timestamps

- PTP Hardware Timestamp
- Software Timestamp

PTP Hardware Timestamp

This setting allows applications that use PTPv2 (Precision Time Protocol) to use hardware generated timestamps to synchronize clocks throughout your network. If this setting is enabled, it takes precedence over the Software Timestamp setting.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "PTP HardwareTimestamp" -DisplayValue "Enabled"
```

Possible values for this setting are:

- Enabled
- Disabled

Software Timestamp

This setting allows applications that use PTPv2 (Precision Time Protocol) to use software generated timestamps to synchronize clocks throughout your network. If the PTP Hardware Timestamp setting is enabled, it takes precedence over this setting.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Software Ti
mestamp"-DisplayValue "RxAll"
```

Possible values for this setting are:

- Disabled
- RxAll
- TxAll
- RxAll & TxAll
- TaggedTx
- RxAll & TaggedTx

Transmit Buffers

This setting 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.

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Transmit Bu
ffers"-DisplayValue "128"
```

Possible values for this setting are:

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

UEFI Network Device Drivers

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.

You can download reference implementations of the PXE and IP based network stack source code at https://www.tianocore.org/.

Loading the UEFI Network Driver

Load the network driver using the UEFI shell **load** command. For example:

```
load e3040e2.efi
```

Configuring the 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 that can be enabled include:

- Windows Deployment Services (WDS)
- Linux network installation (Elilo)
- TFTP file transfers

To enable UEFI PXE services, the following network protocol drivers must be loaded with:

- snp.efi
- bc.efi
- pxedhcp4.efi

These drivers can be loaded from the UEFI **load** shell command, but are often included as part of the UEFI system firmware.

Use the UEFI shell command **drivers** 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

DRV	VERSION	TYPE	CFG	DIAG	#D	#C	DRIVER NAME	IMAGE NAME
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 the UEFI Network Stack for TCP, UDP, and 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
- 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

DRV	VERSION	TYPE	CFG	DIAG	#D	#C	DRIVER NAME	IMAGE NAME
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, use the UEFI **unload** command. The syntax for using the **unload** command is:

```
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. You can access the force speed and duplex menu with the **drvcfg** UEFI shell command:

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

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 (and faster) 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.

The following performs a basic hardware register test:

```
drvdiag -s
```

The following performs an internal loopback transmit and receive test:

```
drvdiag -e
```

VF Loopback Pacing

On Linux* and Microsoft Windows Server*, the Intel® Ethernet 800 Series supports adjusting the loopback rate for a designated port, which allows you to prioritize that port for maximum bandwidth and achieve higher speeds.

- On Linux, use the **devlink** command and the loopback parameter to change this setting.
- On Windows Server, use Ethernet Cmdlets for Intel® Ethernet or Intel® PROSet to change the Loopback setting.

  > **Note:** This feature is only supported on Windows Server 2019 and newer versions.

After changing the loopback setting, the driver will reconfigure all underlying VFs to align to the desired port settings and add more bandwidth to VF-to-VF traffic. Setting this parameter to prioritized enables higher hairpin-bandwidth on related PFs.

> **Note:**
>
> - This configuration applies only for 8x10G and 4x25G adapter cards.
> - Typically you would set some ports to prioritized loopback and then disable loopback on other ports, to allow the driver to utilize spare bandwidth for VF-to-VF traffic.

- Intel recommends using the prioritized loopback setting on a port with minimal network traffic.
- You should first configure loopback on the PF and then configure any other settings, such as VFs/VMs or assigning MAC addresses.

To change this setting with Ethernet Cmdlets for Intel Ethernet or Intel PROSet:

This setting is found in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting with Ethernet Cmdlets for Intel Ethernet, use the Set-IntelEthernetSetting cmdlet. For example:

```
Set-IntelEthernetSetting -Name "<adapter_name>" -DisplayName "Loopback" -DisplayValue "Prioritized"
```

To change this setting in Intel® PROSet for Windows PowerShell* software, use the Set-IntelNetAdapterSetting cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "Loopback" -DisplayValue "Prioritized"
```

Possible values for this setting are:

- Enabled
- Disabled
- Prioritized

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.

See the following subsections for more information.

- Single Root I/O Virtualization (SR-IOV)
- Virtual Machine Queue Offloading
- Using Intel Network Adapters in a Microsoft* Hyper-V* Environment

Note:

- 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.
- In Microsoft Windows Server* 2022 and later, the number of virtual functions (VFs) that the Windows driver advertises to the host may differ from the number of VFs advertised in PCIe configuration space, due to a limitation of hardware resources such as receive queue resources.

Single Root I/O Virtualization (SR-IOV)

SR-IOV Overview

Single Root I/O 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.

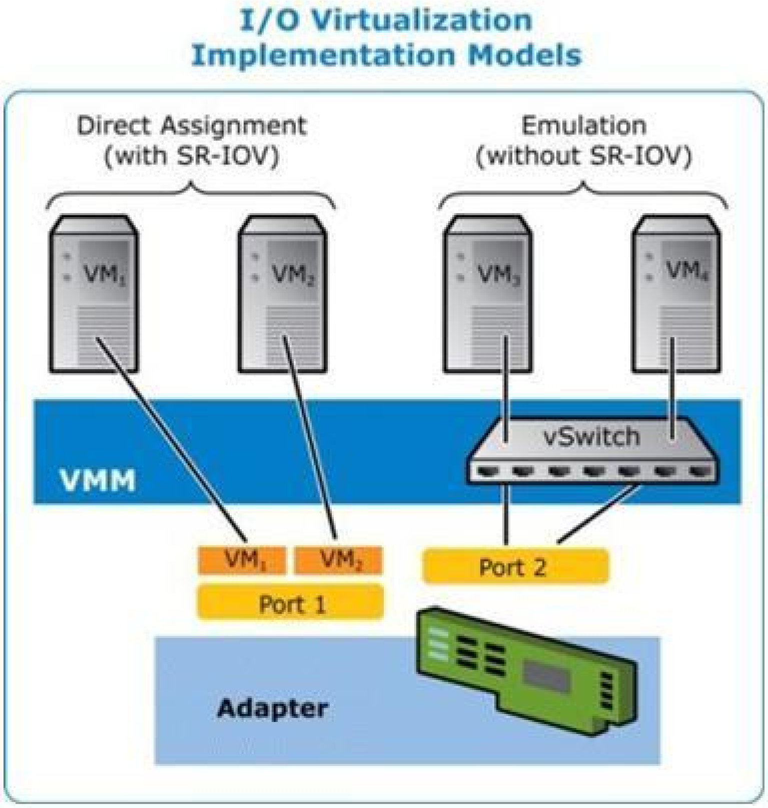
I/O Virtualization Implementation Models
Direct Assignment (with SR-IOV)
Emulation (without SR-IOV)
VM_1
VM_2
VM_3
VM_4
vSwitch
VMM
VM_1
VM_2
Port 1
Port 2
Adapter

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.

Configuring SR-IOV

SR-IOV lets a single network port appear to be several virtual functions in a virtualized environment. If you have an SR-IOV capable device, each port on that device 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. See your operating system documentation for system requirements.

For devices that support it, SR-IOV is enabled in the host partition. Some devices may need to have SR-IOV enabled in a preboot environment.

Note:

- **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 layer two frames, like IEEE 802.3x (link flow control), IEEE 802.1Qbb (priority based flow-control), and others of this type, are not expected and can throttle traffic between the host and the virtual switch, reducing performance. To resolve this issue, and to ensure isolation from unintended traffic streams, configure all SR-IOV enabled ports for VLAN tagging from the administrative interface on the PF. This configuration allows unexpected, and potentially malicious, frames to be dropped.
- SR-IOV must be enabled in the BIOS.
- 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.
- Binding more than two virtual functions (VFs) to a virtual machine (VM) is not recommended. Binding more VFs to a VM may cause system instability.
- SR-IOV is not supported with Intel ANS teams.
- VMWare ESXi* does not support SR-IOV on 1Gbps ports.

Configuring SR-IOV in Windows*

Use Intel® PROSet to change this setting in Windows.

This setting is found on the Advanced tab of the device's Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

To change this setting in Windows PowerShell*, use the `Set-IntelNetAdapterSetting` cmdlet. For example:

```
Set-IntelNetAdapterSetting -Name "<adapter_name>" -DisplayName "SR-IOV" -D
isplayValue "Enabled"
```

Virtual Machine Queue Offloading

Enabling Virtual Machine Queue (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.

Each Intel® Ethernet Adapter has a pool of virtual ports that are split between the various features, such as VMQ Offloading, SR-IOV, and Data Center Bridging (DCB). 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.

Note:

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

Enabling VMQ Offloading in Windows*

For devices that support it, VMQ offloading is enabled in the host partition in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU) or on the Advanced tab of the adapter's Device Manager property sheet, under Virtualization properties.

Virtualization properties also displays the number of virtual ports available for virtual functions, and allows you to set the distribution of available virtual ports 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.

See Adapter Teaming for more information on teams.

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.

Using Intel Network Adapters in a Microsoft* Hyper-V* Environment

When a Hyper-V Virtual NIC (VNIC) interface is created in the host OS, the VNIC takes on the MAC address of the underlying physical NIC (PF, or physical function). 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), 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.

Note:

- 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.

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 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® Advanced Network Services (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 Intel ANS VLANs

Note:

See Adapter Teaming and Virtual LANs (VLANs) for more information on Intel ANS and VLANs.

If you create Intel ANS VLANs in the host OS, and you then create a Hyper-V Virtual NIC interface on an Intel ANS VLAN, then the Virtual NIC interface must have the same VLAN ID as the 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 Intel 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 Microsoft Windows* Device Manager.

Using an Intel 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:

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 host OS.)
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.

> **Note:**
>
> - These steps apply 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.

Command Line for Microsoft Windows Server Core

Microsoft Windows Server* Core does not have a GUI interface. If you want to use an Intel 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.

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 VirtualMachine
LoadBalancing -TeamName "Team1"
```

Virtual LANs (VLANs)

Overview

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. For example:

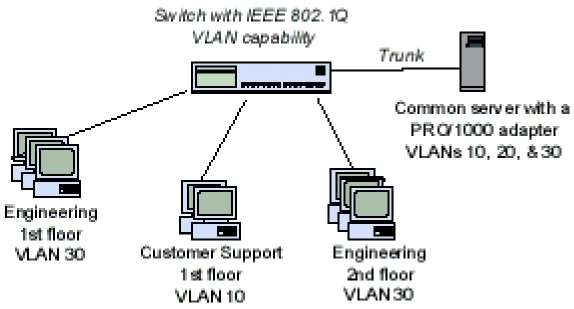

Typically, VLANs consist of co-workers within the same department but in different locations, groups of users running the same network protocol, or a cross-functional team working on a joint project.

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

VLANs and Intel® Advanced Network Services (Intel® ANS)

For more information on Intel ANS, refer to Adapter Teaming.

- Intel ANS is not supported on Microsoft Windows Server* 2016 and later.
- Microsoft Windows* 10 is the last Windows operating system version that supports Intel ANS. Intel ANS is not supported on Microsoft Windows 11 and later.

- 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.
- 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.

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 layer two frames, like IEEE 802.3x (link flow control), IEEE 802.1Qbb (priority based flow-control), and others of this type, are not expected and can throttle traffic between the host and the virtual switch, reducing performance. To resolve this issue, and to ensure isolation from unintended traffic streams, configure all SR-IOV enabled ports for VLAN tagging from the administrative interface on the PF. This configuration allows unexpected, and potentially malicious, frames to be dropped.
- The VF is not aware of the VLAN configuration if you use LBFO to configure VLANs in a Windows guest. VLANs configured using LBFO on a VF driver may result in failure to pass traffic. You must use Windows Hyper-V on the host to configure VLANs on a Windows guest.
- 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.
- Intel ANS VLANs can co-exist with Intel ANS teams (if the adapter supports both). If you do this, the team must be defined first, then you can set up your VLAN.
- You can set up only one untagged VLAN per adapter or team. You must have at least one tagged VLAN before you can set up an untagged VLAN.
- Jumbo Frames are not supported over Intel ANS VLANs under Microsoft Windows 10.

Note:

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

Configuring VLANs in Microsoft Windows

Using Windows PowerShell*

To add a VLAN, use the Add-IntelNetVLAN cmdlet. For example:

```
Add-IntelNetVLAN -ParentName "Name" -VLANID "1"
```

To remove a VLAN, use the Remove-IntelNetVLAN cmdlet. For example:

```
Remove-IntelNetVLAN -ParentName "Name" -VLANID "1"
```

Using Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU)

On the Teaming/VLANs tab, use the VLANs panel.

Using Intel® PROSet for Windows* Device Manager

This setting is found on the VLANs tab of the device's Device Manager property sheet.

Note:

Do not use the Network Connections dialog box to enable or disable VLANs. Otherwise, the VLAN driver may not be correctly enabled or disabled.

Note:

- 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.
- In most environments, a maximum of 64 VLANs per network port or team are supported by Intel ANS.
- Intel 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

This setting 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”

To change this setting in Intel® PROSet:

This setting is found on the Advanced tab of the device’s Device Manager property sheet or in the Adapter Settings panel in Intel® PROSet Adapter Configuration Utility (Intel® PROSet ACU).

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"
```

Possible values for this setting are:

- On
- Off
- Auto Detect

Tools & Apps

This section describes the tools and applications available for Intel® Ethernet devices, adapters, and connections.

- Intel® Network Connection Tools
  - Installing and Uninstalling Intel® Network Connection Tools
  - Intel® Ethernet Flash Firmware Utility
  - Intel® Dynamic Device Personalization Tool
  - Intel® DIAGS
  - Ethernet Port Configuration Tool (EPCT)
  - Intel® Ethernet NVM Update Tool
  - NVM Update Reference Application (NVM URA) for Intel® Ethernet
  - NVM Update Package
- Ethernet Cmdlets for Intel® Ethernet and Intel® PROSet
  - Compatibility Notes for Ethernet Cmdlets for Intel® Ethernet and Intel® PROSet
  - About Ethernet Cmdlets for Intel® Ethernet
    - Installing Ethernet Cmdlets for Intel® Ethernet
    - Diagnostics in Ethernet Cmdlets for Intel® Ethernet
  - About Intel® PROSet
    - Installing Intel® PROSet

    - Configuring Features with Intel® PROSet
    - Diagnostics in Intel® PROSet
  - Configuring Features with Windows PowerShell*
  - Changing Intel® Ethernet Settings Under Windows Server* Core