


Dual Port PCI Express (PCIe x4) Gigabit Ethernet Server Adapter Network Card - Intel® i350 NIC

ST2000SPEXI



*actual product may vary from photos

DE: Bedienungsanleitung - de.startech.com

FR: Guide de l'utilisateur - fr.startech.com

ES: Guía del usuario - es.startech.com

IT: Guida per l'uso - it.startech.com

NL: Gebruiksaanwijzing - nl.startech.com

PT: Guia do usuário - pt.startech.com

For the most up-to-date information, please visit: www.startech.com

FCC Compliance Statement

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Use of Trademarks, Registered Trademarks, and other Protected Names and Symbols

This manual may make reference to trademarks, registered trademarks, and other protected names and/or symbols of third-party companies not related in any way to StarTech.com. Where they occur these references are for illustrative purposes only and do not represent an endorsement of a product or service by StarTech.com, or an endorsement of the product(s) to which this manual applies by the third-party company in question. Regardless of any direct acknowledgement elsewhere in the body of this document, StarTech.com hereby acknowledges that all trademarks, registered trademarks, service marks, and other protected names and/or symbols contained in this manual and related documents are the property of their respective holders.

Table of Contents

Introduction	2
Packaging Contents	3
System Requirements.....	3
Installation	3
Hardware Installation	3
Driver Installation.....	4
Windows	4
Verifying Installation.....	5
Windows	5
Advanced Features (Windows)	6
Link Speed & Duplex	6
Advanced Tab.....	7
Gigabit Master Slave Mode	7
Interrupt Moderation	8
Jumbo Packet.....	8
Large Send Offload	9
Locally Administered Address.....	9
Log Link State Event.....	10
Maximum Number of RSS Queues	10
Packet Priority & VLAN	11
Performance Options	11
Receive Side Scaling	12
TCP/IP Offloading Options	12
Wait for Link.....	13

Power Management.....	13
Teaming	15
ANS Teaming Types	18
Primary and Secondary Adapters	21
Failover and Failback	21
VLANs	22
Specifications.....	23
Technical Support	24
Warranty Information.....	24

Introduction

The ST2000SPEXI PCI Express Dual Port Gigabit Network Card lets you add two gigabit Ethernet ports to a desktop computer through a single, four-lane (x4 or higher) PCI Express slot. Each port is fully compatible with IEEE 802.3/u/ab standards, for 10/100/1000 Mbps auto-negotiation.

This 2-port GbE network card is based on the high-performance Intel® I350 chipset, which offers several advanced driver options including IEEE 802.3ad link aggregation / teaming support, enabling you to add additional bandwidth to your system. Support for Jumbo Frames, full-duplex operation and 802.1q VLAN tagging also aid in creating an efficient network configuration.

Also useful for virtualization applications or setting up a dual-homed host configuration, this dual-port network card lets you accommodate several complex network scenarios, to meet your business needs. For added versatility, the controller card is equipped with a standard profile bracket and includes a low-profile/half-height bracket for installation in small form-factor computers.

Packaging Contents

- 1 x 2 Port Network Card
- 1 x Low Profile Bracket
- 1 x Driver CD
- 1 x Instruction Manual

System Requirements

- Available PCI Express x4 or higher (x8, x16) slot
- Windows® 8.1 (32/64bit), 8 (32/64bit), 7 (32/64), Vista(32/64), XP(32/64), Windows® Server 2012, 2008 R2, 2003(32/64), Linux 2.4.x/ 2.6.x

Installation

Hardware Installation

WARNING! PCI Express cards, like all computer equipment, can be severely damaged by static electricity. Be sure that you are properly grounded before opening your computer case or touching your PCI Express card. StarTech.com recommends that you wear an anti-static strap when installing any computer component. If an anti-static strap is unavailable, discharge yourself of any static electricity build-up by touching a large grounded metal surface (such as the computer case) for several seconds. Also be careful to handle the card by its edges and not the gold connectors.

1. Turn your computer off and any peripherals connected to the computer (i.e. Printers, external hard drives, etc.). Unplug the power cable from the rear of the power supply on the back of the computer and disconnect all peripheral devices.
2. Remove the cover from the computer case. See documentation for your computer system for details.
3. Locate an open PCI Express slot and remove the metal cover plate on the rear of the computer case (Refer to documentation for your computer system for details.). Note that this card will work in PCI Express slots of additional lanes (i.e. x8 or x16 slots).
4. Insert the card into the open PCI Express slot and fasten the bracket to the rear of the case.

Note: If installing the card into a low profile desktop system, replacing the preinstalled standard profile bracket with the included low profile (half height) installation bracket may be necessary.

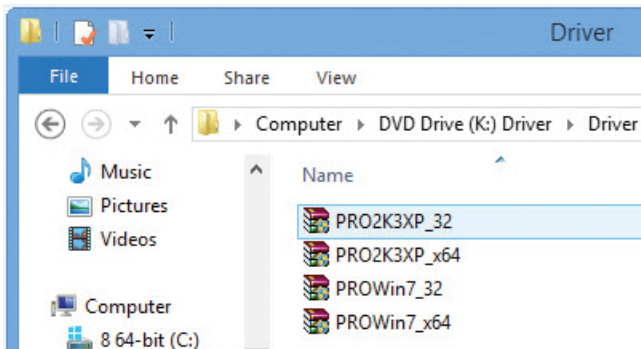
5. Place the cover back onto the computer case.
6. Insert the power cable into the socket on the power supply and reconnect all other connectors removed in Step 1.

Driver Installation


Windows

Note: The card may auto-install using native drivers, however it is recommended to use the procedure below to update to the latest version and have access to advanced features.

1. Upon starting Windows, if the "Found New Hardware" wizard appears on the screen, cancel/close the window and insert the included Driver CD into the computer's CD/DVD drive.
2. AutoPlay should automatically launch and display all of the folders on the CD or prompt you to view files. If AutoPlay is disabled on your system, open a **Computer** window and navigate to your CD/DVD drive.
3. Choose the appropriate driver for your version of windows and double-click to run the installer.




4. Choose the driver options you would like installed and follow the on-screen instructions to complete the installation.



Verifying Installation

Windows

1. Open the Device Manager by right-clicking on **Computer**, and then select **Manage**. In the new Computer Management window, select **Device Manager** from the left window panel (For Windows 8, open the **Control Panel** and select **Device Manager**).
2. Expand the “Network Adapters” section. On a successful install, you should following devices in the list with no exclamation points or question marks.



Advanced Features (Windows)

Link Speed & Duplex

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

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

Manually setting Speed & Duplex

By default, auto-negotiation is enabled. Change this setting only to match your link partner.

1. Open the Device Manager and expand the Network Adapters section.
2. Right-click the Intel i350 Gigabit Network Connection you wish to configure and select Properties.
3. Click the **Link Speed** tab.
4. Select the appropriate speed and duplex from the Speed and Duplex pull down menu.
5. Click **OK**.

Gigabit Network Device Considerations

Per the IEEE specification, gigabit speed is available only in full duplex.

The settings available when auto-negotiation is disabled are:

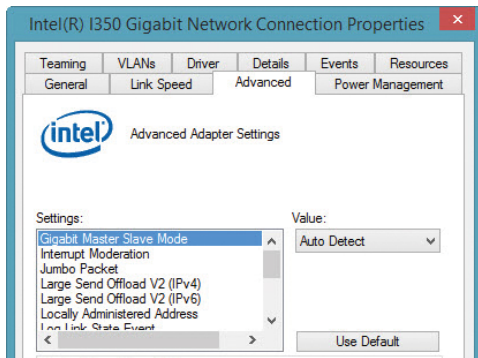
- **10 Mbps or 100 Mbps Full duplex** (requires a full duplex capable link partner set to full duplex). The adapter can send and receive packets at the same time. You must set this mode manually.
- **10 Mbps or 100 Mbps Half duplex** (requires a link partner set to half duplex). The adapter performs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

- **Auto-Negotiation 1000 Mbps.** The adapter only advertises gigabit speed at full duplex.

Advanced Tab

All of the settings explained in the following sections can be accessed / modified from the Advanced tab in the device properties:

1. Open the Device Manager and expand the **Network Adapters** section.
2. Right-click the Intel i350 Gigabit Network Connection you wish to configure and select **Properties**.
3. Click the **Advanced** tab.
4. Make your desired modifications.
5. Click OK to apply the changes.



Gigabit Master Slave Mode

Decides whether the adapter or link partner, is designated as the master. The other device is designated as the slave. Changing the setting can improve link quality with certain link partners.

NOTE: Some multi-port devices can be forced to master mode. If forced to master mode, the device can either disconnect or downshift to a 100-Mbps link. This issue can also occur in a Forced Slave to Forced Slave mode. Changing this setting can cause a momentary loss of connectivity.

Interrupt Moderation

Sets the Interrupt Throttle Rate (ITR), the rate at which the controller moderates interrupts.

The default setting is optimized for common configurations. Changing this setting can improve network performance on certain network and system configurations.

Options:

Adaptive (ITR = -1, no interrupts/sec, it is dynamically changed by the driver)

Off (ITR = 0, no limit)

Minimal (ITR = 200)

Low (ITR = 400)

Medium (ITR = 950)

High (ITR = 2000)

Extreme (ITR = 3600)

When an event occurs, the adapter generates an interrupt, which allows 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 system performance.

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

Jumbo Packet

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

NOTES:

- All equipment on the network must also support the larger frame size. When setting up Jumbo Frames on other network devices, be aware that different network devices calculate Jumbo Frame size differently. Some devices include the header information in the frame size while others do not. Intel adapters do not include header information in the frame size. When configuring Jumbo Frames on a switch, set the frame size four bytes higher for CRC, plus four bytes if you are using VLANs or QoS packet tagging.
- Jumbo Frames are supported at 1000 Mbps only. Using Jumbo Frames at 10 or 100 Mbps is not supported and may result in poor performance or loss of link.

Restrictions:

- Jumbo Frames can be implemented simultaneously with VLANs and teaming, however 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. To configure Jumbo Frames at the switch, consult your network administrator or switch user's guide.
- When standard sized Ethernet frames (64 to 1518 bytes) are used, there is no benefit to configuring Jumbo Frames.
- The Jumbo Frames setting on the switch must be set to at least 8 bytes larger than the adapter setting for Microsoft Windows operating systems, and at least 22 bytes larger for all other operating systems.

Large Send Offload

Enables the adapter to offload the task of segmenting TCP messages into valid Ethernet frames.

Because the adapter hardware can complete data segmentation much faster than operating system software, this feature can improve transmission performance. The adapter also uses fewer CPU resources.

Large Send Offload can be configured separately for IPv4 and IPv6.

Locally Administered Address

The Locally Administered Address feature allows you to override a device's MAC Address and assign an address of your choice. The address is a 12-digit hexadecimal number in the range: 0000 0000 0001 – FEFF FFFF FFFF.

CAUTION: Make sure no other systems on the network use this address.

NOTES:

- 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.)
- To restore the default MAC address, click **Use Default**.
- You cannot change the address if the adapter is part of a team.
- Changing this setting may cause a momentary loss of connectivity.

The LAA setting is disabled and cannot be changed if one or more of the following is true:

- The device is configured to iSCSI boot primary or secondary.
- The device or team is bound to the Microsoft Virtual Switch Protocol.

Log Link State Event

Enables the logging of the following link state changes to the system event log:

LINK_UP_CHANGE - When this message is displayed, the link is up.

LINK_DOWN_CHANGE - When this message is displayed, the link is down. To investigate this issue, click the **Link Speed** tab and run diagnostics.

LINK_DUPLEX_MISMATCH - When this message is displayed, there is a mismatch in duplex between the adapter and the link partner. To investigate this issue, click the **Link Speed** tab and change the speed and duplex settings appropriately.

Maximum Number of RSS Queues

Configures the number of Receive Side Scaling (RSS) queues:

- One queue is used when low CPU utilization is required.
- Two queues are used when good throughput and low CPU utilization are required.
- Four or more queues are used for applications that demand high transaction rates such as web server based applications. With this setting, the CPU utilization may be higher.

NOTES:

- RSS is not supported on some adapters configured to use Virtual Machine Queues (VMQ). For these adapters VMQ takes precedence over RSS. RSS is disabled.
- Using eight or more queues on Microsoft Windows Server 2008 requires the system to reboot.
- Changing this setting may cause a momentary loss of connectivity.

Packet Priority & VLAN

Enables sending and receiving of IEEE 802.3ac tagged frames, which include:

802.1p QoS (Quality of Service) tags for priority-tagged packets

802.1Q tags for VLANs

Intel® PROSet automatically enables 802.1p tagging when you set up a VLAN. You cannot disable QoS Packet Tagging on a VLAN because tagging is required for VLANs.

When this feature is enabled, tagged packets use the queue settings defined by the operating system Priority Level Definition.

When this setting is disabled, the adapter cannot tag outgoing packets with 802.1p/802.1Q tags.

Performance Options

Receive Descriptors or Receive Buffers

Sets the number of buffers used by the driver when copying data to the protocol memory. Increasing this value can enhance the receive performance, but also consumes system memory.

Receive Descriptors are data segments that enable the adapter to allocate received packets to memory. Each received packet requires one Receive Descriptor, and each descriptor uses 2 KB of memory.

Transmit Descriptors or Transmit Buffers

Defines the number of Transmit Descriptors, 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 Descriptors.

You might choose to increase the number of Transmit Descriptors if you notice a problem with transmit performance. Increasing the number of Transmit Descriptors can enhance transmit performance. But, Transmit Descriptors consume system memory. If transmit performance is not an issue, use the default setting.

Receive Side Scaling

Enables Receive Side Scaling (RSS). RSS balances receive traffic across multiple CPUs or CPU cores. This setting has no effect if your system has only one processing unit.

NOTES:

- RSS must be enabled for Intel® I/O Acceleration Technology to function under Microsoft Windows Server 2003.
- RSS is not supported on some adapters configured to use Virtual Machine Queues (VMQ). For these adapters VMQ takes precedence over RSS. RSS is disabled.
- Changing this setting can cause a momentary loss of connectivity.

Teaming Notes:

- If RSS is not enabled for all adapters in a team, RSS is disabled for the team.
- If an adapter that does not support RSS is added to a team, RSS is disabled for the team.
- If a non-Intel adapter is added to a team, its RSS settings must match the Intel adapters in the team.
- Not all adapters support all RSS queue settings

TCP/IP Offloading Options

IPv4 Checksum Offload

Enables the adapter to verify the IP checksum on received packets and compute checksum on transmitted packets. Enabling this feature can improve IP performance and reduce CPU utilization.

NOTE: Changing this setting may cause a momentary loss of connectivity.

TCP Checksum Offload

Enables the adapter to verify the TCP checksum on received packets and compute checksum on transmitted packets. Enabling this feature can improve TCP performance and reduce CPU utilization.

This setting can be configured separately for IPv4 and IPv6.

NOTE: Changing this setting may cause a momentary loss of connectivity.

Transmit IP Checksum

Allows the adapter to compute the IP checksum of transmitted packets. This feature can improve IP transmit performance and reduce CPU utilization. With Offloading disabled, the operating system verifies the IP checksum.

UDP Checksum Offload

Enables the adapter to verify the UDP checksum on received packets and compute checksum on transmitted packets. Enabling this feature can improve UDP performance and reduce CPU utilization.

This setting can be configured separately for IPv4 and IPv6.

Wait for Link

Decides if the driver waits for Auto Negotiation to be successful before reporting the link state.

Options

Off: Driver does not wait for Auto Negotiation

On: Driver does wait for Auto Negotiation. If the speed is not set to Auto Negotiation, the driver waits for a short time for link to complete and then reports the link state.

Auto Detect: Automatically set to On or Off depending on speed and adapter type when the driver is installed.

- Off for copper Intel PRO/1000 adapters with a speed of Auto
- On for copper Intel PRO/1000 adapters with a forced speed and duplex
- On for fiber Intel PRO/1000 adapters with a speed of Auto

Power Management

The Intel® PROSet Power Management tab replaces the standard Microsoft Windows Power Management tab in Device Manager. It includes the Power Saver and Wake on LAN (WoL) options that were previously included on the Advanced tab. The standard Windows power management functionality is incorporated on the Intel PROSet tab. The Intel PROSet Power Management tab includes several settings that control the adapter's power consumption.

Energy Efficient Ethernet

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

NOTES:

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

Remote Wake-Up and Wake on LAN


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

NOTE:

- In some scenarios, the "Wait for Link" option (located on the Advanced Tab) needs to be set to On for Wake on LAN to function properly. If you are experiencing any difficulty getting WoL to work, please adjust this setting and try again.

Configuring Power Management Settings

1. Open the Device Manager and expand the Network Adapters section.
2. Right-click the Intel i350 Gigabit Network Connection you wish to configure and select Properties.
3. Click the Power Management tab.
4. Set your desired power management options.
5. Click OK to apply the changes.



Teaming

Advanced Network Services (ANS) Teaming, a feature of the Advanced Network Services component, lets you take advantage of multiple network ports / adapters in a system by grouping them together. ANS teaming can use features like fault tolerance and load balancing to increase throughput and reliability. Before you can set up ANS teaming in Windows, you must install Intel® PROSet software.

IMPORTANT:

- 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 Intel PROSet's **Advanced** tab. Pay particular attention when using different adapter models or adapter versions, as adapter capabilities vary.
- If team members implement advanced features differently, failover and team functionality will be affected. To avoid team implementation issues, create teams that use similar adapter types and models.
- Reload the team after adding an adapter or changing any advanced features. One way to reload the team is to select a new preferred primary adapter. Although there will be a temporary loss of network connectivity as the team reconfigures, the team will maintain its network addressing schema.

NOTES:

- Be sure to use the latest available drivers.
- 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.
- 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.
- Intel adapters that do not support Intel PROSet may still be included in a team. However, they are restricted in the same way non-Intel adapters are. See **Multi-Vendor Teaming** below for more information.
- To assure a common feature set, some advanced features, including hardware offloading, are automatically disabled when an adapter that does not support Intel PROSet is added to a team.
- Hot Plug operations in a Multi-Vendor Team may cause system instability. We recommend that you restart the system or reload the team after performing Hot Plug operations with a Multi-Vendor Team.
- Spanning tree protocol (STP) should be disabled on switch ports connected to teamed adapters in order to prevent data loss when the primary adapter is returned to service (failback). Alternatively, an activation delay may be configured on the adapters to prevent data loss when spanning tree is used. Set the Activation Delay on the advanced tab of team properties.
- ANS teaming of Virtual Function devices inside a Windows 2008 R2 guest running on an open source hypervisor is supported.
- Fiber Channel over Ethernet/Data Center Bridging will be automatically disabled when an adapter is added to a team with non-FCoE/DCB capable adapters.

Conditions that may prevent you from teaming a device

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

- The operating system does not support the desired team type.
- The device does not support the desired team type or does not support teaming at all.
- The devices you want to team together use different driver versions.
- TOE (TCP Offload Engine) enabled devices cannot be added to an ANS team and will not appear in the list of available adapters.
- You can add Intel® Active Management Technology (Intel® AMT) enabled devices to Adapter Fault Tolerance (AFT), Switch Fault Tolerance (SFT), and Adaptive Load Balancing (ALB) teams. All other team types are not supported. The Intel AMT enabled device must be designated as the primary adapter for the team.
- The device's MAC address is overridden by the Locally Administered Address advanced setting.
- The device has a virtual NIC bound to it.
- The device is part of a Microsoft Load Balancing and Failover (LBFO) team.

Microsoft Load Balancing and Failover (LBFO) teams

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

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

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.

ANS Teaming Types

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

Switch Fault Tolerance (SFT) - provides failover between two adapters connected to separate switches. Switch Fault Tolerance supports two adapters per team. Spanning Tree Protocol (STP) must be enabled on the switch when you create a SFT team. During normal operation, the secondary adapter is in standby. When SFT teams are created, the Activation Delay is automatically set to 60 seconds. This teaming type works with any switch or hub. All team members must be connected to the same subnet. In SFT teams, each adapter in the team can operate at a different speed than the other.

- 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.
- Switch Fault Tolerance can detect failures when they occur on either teamed adapter, on either cable connecting the teamed adapter to its switch, or on switch ports connected to the adapters, if link is lost

Adaptive Load Balancing (ALB) - uses software to balance routable traffic among a team of two to eight adapters or LOMs (the team must include at least one server adapter) connected to the same subnet. The software 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 also provide the benefits of fault tolerance.

NOTES:

- ALB does not load balance non-routed protocols such as NetBEUI and some IPX* traffic.
- You can create an ALB team with mixed speed adapters. The load is balanced according to the lowest common denominator of adapter capabilities and the bandwidth of the channel.

- On Windows systems, Receive Load Balancing is enabled by default.
- Receive Load Balancing is not supported on Microsoft Hyper-V.

Virtual Machine Load Balancing (VMLB) - provides transmit and receive traffic load balancing across Virtual Machines bound to the team interface, as well as fault tolerance in the event of switch port, cable, or adapter failure. This teaming type works with any switch. The driver analyzes the transmit and receive load on each member adapter and balances the traffic across member adapters. In a VMLB team, each Virtual Machine is associated with one team member for its TX and RX traffic. If only one virtual NIC is bound to the team, or if Hyper-V is removed, then the VMLB team will act like an AFT team.

NOTES:

- VMLB does not load balance non-routed protocols such as NetBEUI and some IPX* traffic.
- VMLB supports from two to eight adapter ports per team.
- You can create an 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) - provides increased transmission and reception throughput in a team of two to eight adapters. This type also includes adapter fault tolerance and load balancing (only routed protocols). This teaming type requires a switch with Intel Link Aggregation, Cisco* FEC or GEC, or IEEE 802.3ad Static Link Aggregation capability.

All adapters in a Link Aggregation team running in static mode must run at the same speed and must be connected to a Static Link Aggregation capable switch. If the speed capabilities of adapters in a Static Link Aggregation team are different, the speed of the team is dependent on the lowest common denominator. **Note:** You cannot use an Intel AMT enabled adapter in an SLA team

IEEE 802.3ad Dynamic Link Aggregation - creates one or more teams using Dynamic Link Aggregation with mixed-speed adapters. Like the Static Link Aggregation teams, Dynamic 802.3ad teams increase transmission and reception throughput and provide fault tolerance. This teaming type requires a switch that fully supports the IEEE 802.3ad standard.

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

NOTES:

- Once you choose an aggregator, it remains in force until all adapters in that aggregator 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 copper or fiber-based adapters in a system.

Before you begin

- 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. If needed, change your speed or duplex setting before you link the adapter to the switch. 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 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 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) - adds the capability to include adapters from selected other vendors in a team. If you are using a Windows Server based computer, all 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 that supports Intel PROSet in the team. That adapter or connection must be designated as the primary adapter.
- A multi-vendor team can be created for any team type.

- All members in a MVT must operate on a common feature set (lowest common denominator).
- Manually verify that the frame setting for all adapters in the team is the same.

Verify that the RSS settings for all adapters in the team are the same.

Primary and Secondary Adapters

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. The adapter will, however, rejoin the team as a non-primary.
- **Preferred Primary/Secondary adapters:** You can specify a preferred adapter in Intel PROSet. Under normal conditions, the Primary adapter handles all non-TCP/IP traffic. The Secondary adapter will receive fallback traffic if the primary fails. If the Preferred Primary adapter fails, but is later restored to an active status, control is automatically switched back to the Preferred Primary adapter.

To specify a preferred primary or secondary adapter

1. From Device Manager, open the properties of a team.
2. Click the **Settings** tab.
3. Click the **Modify Team** button.
4. Select the adapter you want to be the primary adapter and click the **Set Primary** button.

The adapter's preferred setting appears in the Priority column.

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).

VLANs

VLAN (Virtual Local Area Network) refers to a collection of devices that communicate as if they were on the same physical LAN. 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.

NOTES:

- When using IEEE 802 VLANs, settings must match between the switch and those adapters using the VLANs.
- To set up IEEE VLAN membership (multiple VLANs), the adapter must be attached to a switch with IEEE 802.1Q VLAN capability.
- A maximum of 64 VLANs per network port or team are supported by Intel software.
- VLANs can co-exist with teaming. If 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.
- Intel ANS 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 VLAN. You should not add a port that is already part of an Intel ANS VLAN to an LBFO team, as this may cause system instability.

Configuring VLANs

1. Open the Device Manager and expand the Network Adapters section.
2. Right-click the Intel i350 Gigabit Network Connection you wish to configure and select Properties.
3. Click the VLANs tab.
4. Add, Remove or Modify your VLAN settings as applicable to your network.
5. Click OK to apply the changes.

Specifications

Number of Ports	2
Bus Type	PCI Express
Card Type	Standard Profile (LP bracket incl.)
Port Style	Integrated on Card
Chipset	Intel – i350
Industry Standards	IEEE 802.3, IEEE 802.3u, IEEE 802.3ab PCI Express Base Specification 2.1
Connector(s)	1 - PCI Express x4 Male 2 - RJ-45 Female
Maximum Data Transfer Rate	2000 Mbps (full-duplex), 1000 Mbps (half-duplex)
Compatible Networks	10/100/1000 Mbps
Full Duplex Support	Yes
Auto MDIX	Yes
Supported Protocols	IEEE 802.3ad (link aggregation), IEEE 1588 (time sync) / 802.1AS, IEEE 802.1q (VLAN tagging)
Operating Temperature	0°C to 50°C (32°F to 122°F)
Storage Temperature	-20°C to 60°C (-4°F to 140°F)
Humidity	5~85% RH
Dimensions	67 x 19 x 121 mm
Weight	80 g
Compatible Operating Systems	Windows® 8.1 (32/64bit), 8 (32/64), 7 (32/64), Vista(32/64), XP(32/64) Windows® Server 2012, 2008 R2, 2003(32/64)

Technical Support

StarTech.com's lifetime technical support is an integral part of our commitment to provide industry-leading solutions. If you ever need help with your product, visit www.startech.com/support and access our comprehensive selection of online tools, documentation, and downloads.

For the latest drivers/software, please visit www.startech.com/downloads

Warranty Information

This product is backed by a two year warranty.

In addition, StarTech.com warrants its products against defects in materials and workmanship for the periods noted, following the initial date of purchase. During this period, the products may be returned for repair, or replacement with equivalent products at our discretion. The warranty covers parts and labor costs only. StarTech.com does not warrant its products from defects or damages arising from misuse, abuse, alteration, or normal wear and tear.

Limitation of Liability

In no event shall the liability of StarTech.com Ltd. and StarTech.com USA LLP (or their officers, directors, employees or agents) for any damages (whether direct or indirect, special, punitive, incidental, consequential, or otherwise), loss of profits, loss of business, or any pecuniary loss, arising out of or related to the use of the product exceed the actual price paid for the product. Some states do not allow the exclusion or limitation of incidental or consequential damages. If such laws apply, the limitations or exclusions contained in this statement may not apply to you.

Hard-to-find made easy. At StarTech.com, that isn't a slogan. It's a promise.

StarTech.com is your one-stop source for every connectivity part you need. From the latest technology to legacy products — and all the parts that bridge the old and new — we can help you find the parts that connect your solutions.

We make it easy to locate the parts, and we quickly deliver them wherever they need to go. Just talk to one of our tech advisors or visit our website. You'll be connected to the products you need in no time.

Visit www.startech.com for complete information on all StarTech.com products and to access exclusive resources and time-saving tools.

StarTech.com is an ISO 9001 Registered manufacturer of connectivity and technology parts. StarTech.com was founded in 1985 and has operations in the United States, Canada, the United Kingdom and Taiwan servicing a worldwide market.