White Paper
Network Connectivity

Intel® Advanced Network Services Software Increases Network Reliability, Resilience and Bandwidth

Adapter teaming is a long-proven method for increasing network reliability, resilience and bandwidth scalability. In 1996, Intel Corporation began formal support of adapter teaming techniques with the introduction of Intel® Adapter Fault Tolerance (Intel® AFT) software. Since then, Intel® Advanced Network Services (Intel® ANS) has continued evolving to include new teaming technologies and techniques. Today, Intel ANS includes support for teaming up to eight adapters in Adapter Fault Tolerance, Switch Fault Tolerance, Adaptive Load Balancing for both transmit and receive, Static Link Aggregation and IEEE 802.3ad Dynamic Link Aggregation. Additionally, Intel ANS provides support for teaming server adapters from mixed vendors and includes support for 802.1q VLANs, making Intel ANS one of the most capable and comprehensive tools for supporting server adapter teaming.

Contents

Executive Summary

Network connection fault-tolerance and bandwidth-scaling technologies have been around for many years, dating back to the days of mini- and mainframe computers. Back then, information technology (IT) departments would traditionally run multiple network cables through conduits in ceilings and floors to assure user connectivity to mission-critical applications. Those basic concepts of assuring user connectivity continue today, but with significantly expanded and enhanced capabilities and techniques.

Beginning in 1996, Intel carried the concept forward by introducing fault-tolerance functionality for Intel® architecture 32-bit Local Area Network (LAN) server adapters. Since then, more features have been added to further enhance reliability and to accommodate migration of applications used by and deployed to desktop and laptop users. This includes development and deployment of Intel Advanced Network Services (Intel ANS) software, which is bundled with the drivers for all Intel® PRO Server Adapters. Intel ANS provides a means of configuring and managing network adapter teaming for increased network reliability, resilience and bandwidth, and it can run under a variety of operating systems, including various versions of Microsoft Windows* and Linux.*

This white paper discusses Intel ANS software functionality in a LAN, the protocols supported, and the different load-balancing algorithms used in server adapter teaming. The discussion also includes Network Interface Card (NIC) Native and Multi-Vendor Teaming (MVT) features supported by Intel PRO Server Adapters or NICs and LAN on Motherboard (LOM) connections.

Adapter Teaming Features Supported by Intel® ANS

Intel ANS is implemented as an intermediate driver that resides in the server's interface driver stack. As shown in Figure 1, the Intel ANS location in the stack hierarchy is below the Network Protocol Stack TCP/IP level and the Network Driver Interface Specifications (NDIS) level. Within this intermediate driver level, Intel ANS implements the device teaming features. The supported teaming features include fault tolerance, load balancing, link aggregation, and Virtual LAN (VLAN) tagging at the Data Link Layer for Ethernet PCI, PCI-X, and PCI Express* devices. Additionally, the Intel ANS intermediate driver supports adapter teaming on various Operating Systems (OS), including Microsoft Windows Server 2003, Windows 2000, Windows XP, and Linux.

Adapter teaming consists of grouping a number of physical network adapters into virtual adapters that provide the fault-tolerance and load-balancing functions of teaming. A team can consist of two to eight Ethernet interfaces (NICs or LOM), and, depending on the teaming mode, one or more interfaces can be active. Those healthy interfaces in a team that are not active are in standby mode. All of the active interfaces in a team participate in load-balancing operations by sending and receiving a portion of the total network traffic. Standby interfaces will take over the network traffic in the event that the active interfaces lose their link.

Additionally, multiple 802.1q tagged VLANs and a single untagged VLAN can be configured on an Intel team or on individual adapters. However, to use untagged VLANs, there must be at least one tagged VLAN present on the team. Currently, only teams consisting entirely of Intel® adapters support the addition of VLANs. When VLANs are added, a virtual adapter is created for each VLAN. The benefit of adding VLANs is that they extend the functionality of the team by allowing the team to reside in multiple sub-networks.

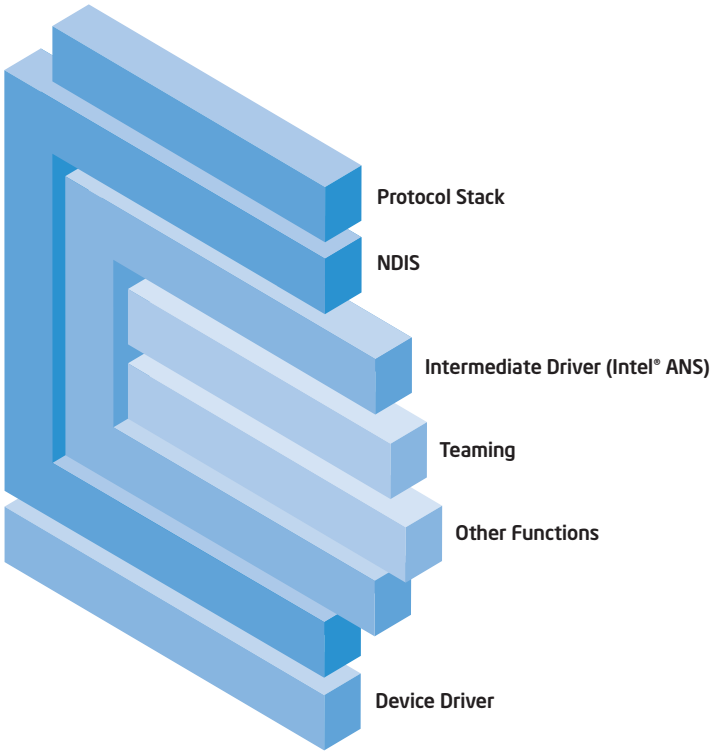

Figure 1. Intel® Advanced Network Services (Intel® ANS) software is implemented as an intermediate driver within the server's driver stack.

In all cases, whether VLANs are included or not, a team must include at least one Intel PRO Server Adapter or an Intel® LOM network connection in order for Intel ANS to be installed and available for adapter teaming. When Intel ANS is installed, the teaming features described below are available.

- **Adapter Fault Tolerance (AFT)** provides automatic redundancy for the server's network connection. If the primary adapter fails, the secondary adapter takes over. AFT supports two to eight adapters per team. This teaming mode works with any hub or switch, and all team members must be connected to the same network. (Note: While a single-interface team can be formed, at least two interfaces must be teamed in order to utilize the fault tolerance feature.)

AFT is inherent in all teaming modes. This mechanism is present in and is the basis for each of the following teaming technologies.

- **Switch Fault Tolerance (SFT)** provides a failover relationship between two adapters when each adapter is connected to a separate switch (see Figure 2). SFT supports two adapters per team and incorporates AFT. It is recommended that Spanning Tree Protocol (STP) be enabled on the network when creating a team in SFT mode; however, STP should be disabled on the switch ports that the teamed server adapters are connected to. SFT is available on computers running Windows 2000 Server, Windows 2003 Server, Windows XP, or Linux. This teaming mode works with any switch.

When implementing Switch Fault Tolerance teaming, use of switches supporting IEEE 802.1w STP is recommended to allow faster failover. Also, only two server adapters can be used in this scenario, and there cannot be more than one adapter in standby mode for proper SFT operation.

- **Adaptive Load Balancing (ALB)** provides load balancing of transmit traffic and fault tolerance. Receive Load Balancing (RLB) is an advanced feature that is enabled by default on ALB teams to provide load balancing of receive traffic. RLB can be disabled or re-enabled in the advanced settings for the team. The ALB/RLB teaming mode works with any switch.
- **Static Link Aggregation (SLA)** supports the Intel® Link Aggregation (LA), Fast EtherChannel* (FEC), Gigabit EtherChannel (GEC), and IEEE 802.3ad* static-mode link-aggregation standards. SLA provides increased transmission and reception throughput in a team comprised of two to eight adapter ports. All ports in an SLA team must be operating at the same Ethernet speed, either Fast Ethernet (100 Mbps) or Gigabit Ethernet. The SLA mode incorporates the AFT and ALB modes. In SLA, if one adapter loses link, the remaining adapters in the team continue to balance the network traffic across the remaining network connections.

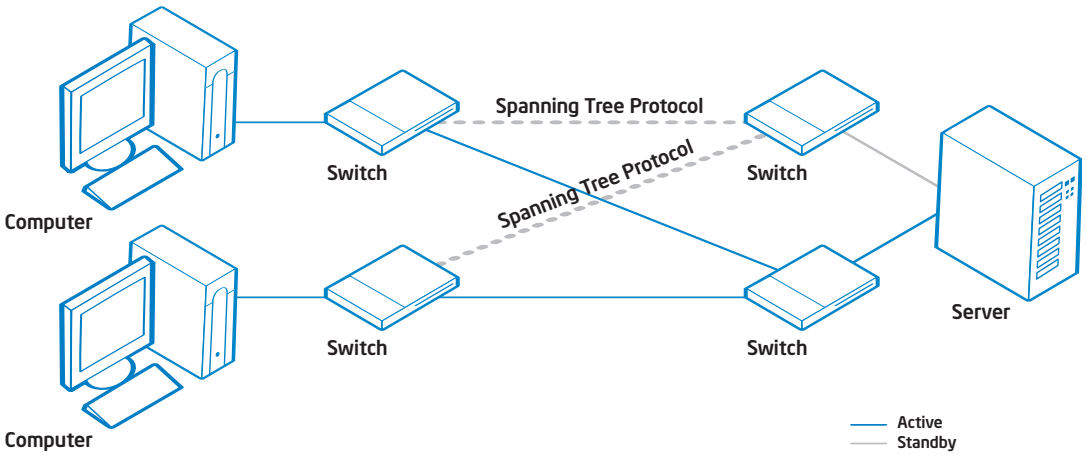

Figure 2. Adapter teaming for Switch Fault Tolerance with Spanning Tree Protocol enabled.

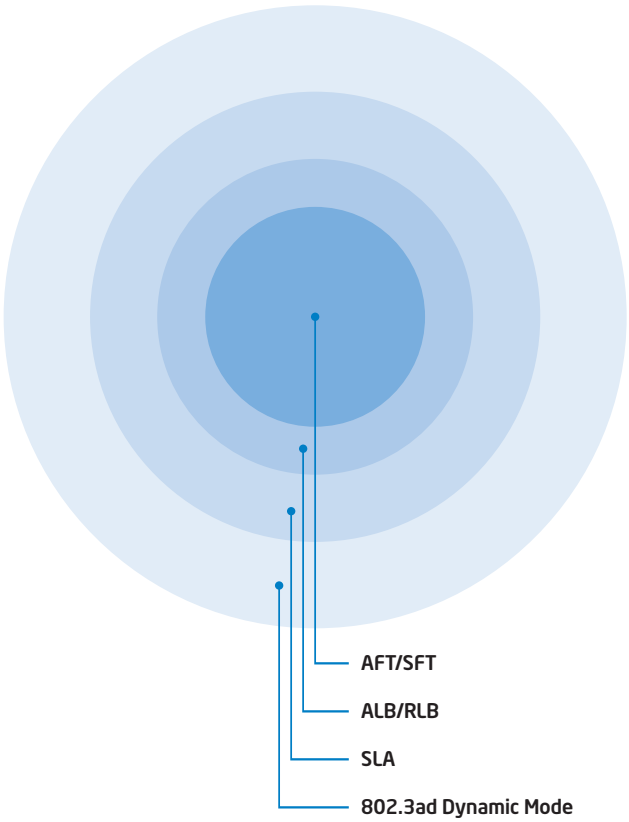

Figure 3. Intel® ANS teaming features build on each other.

All adapters in an SLA team share a virtual Media Access Control (MAC) address derived from the Ethernet address of one of the adapters. Transmit load balancing is based on source/destination address pairs. For this reason, transmit load balancing only occurs when there is more than one destination address. Receive load balancing is determined strictly by the switch that the team is connected to. The switch must support 802.3ad static mode, Port Trunking, Load Sharing or EtherChannels for SLA teams to function properly.

- **IEEE 802.3ad: Dynamic Mode** provides dynamic link aggregation, allowing creation of one or more channel groups using same-speed or mixed-speed server adapters. Using same-speed adapters results in the formation of one channel group, and using mixed-speed adapters results in multiple channel groups. When a channel group in a team fails, the other channel group is activated and provides fault tolerance and link aggregation for the team. A secondary channel group will not fail back to the primary channel unless all members of the secondary group lose link.

Like SLA, 802.3ad dynamic link aggregation increases transmission and reception throughput and provides fault tolerance. The dynamic mode requires use of a switch that fully supports the IEEE 802.3ad standard.

Figure 3 illustrates the inclusiveness of each teaming mode. For example, AFT is the core technology found in ALB. This means if a link goes down in an ALB team, the remaining adapters in the team balance the traffic from the server. Static Link Aggregation includes the features of and expands on ALB and AFT, where traffic is balanced both on transmit and on receive across all network connections, and it includes fault tolerance. The 802.3ad Dynamic Mode is all inclusive, building off of SLA to provide link aggregation for the team and fault tolerance for each link-aggregation channel in the team.

Detecting Network Connection Failures

Having described the different types of available teams, the rest of this paper focuses on how these technologies work in practice.

Inherent in all teaming technologies is network-connection fault tolerance or AFT. In order for AFT to work, at least two network connections must be present and have links in the same network. These two network connections indicate to the system whether or not they have network connectivity using several different mechanisms. Briefly, these mechanisms are as follows:

- **Probe Packets**—the teamed adapters send and receive probe packets to determine the presence or state of each adapter in the team.
- **Activity-Based Tolerance**—when probe packets are not used, or do not reach their destination, selection of a Primary (or active) Adapter is made by sensing activity on the link.
- **Link-Based Tolerance**—when neither probe packets nor activity-based tolerance are available or successful, link-based tolerance is used to check the links of the teamed adapters to select a Primary Adapter.

Detection by Probe Packets

The team probing mechanism is one of three methods used to assess the health and accessibility of individual adapters in a team. Each adapter in the team sends one of two types of user-configurable probes, either broadcast probes or multicast probes, onto the network segment where the team resides. Each member adapter uses two flags—Send and Receive—to track the status of sent and received probe packets. When an adapter sends a probe packet onto the network, it sets the two flags to Pending state. When the probe packet is received by a member of the same team, the receiving team member sets its receive flag to ReceiveComplete and also sets the sending adapter's send flag to SendComplete. Failure in this received/sent flagging sequence indicates a potential adapter or link problem; however, Intel ANS will attempt to further verify or resolve the problem through probe retries or a probe burst mechanism.

When an adapter initially fails to send or receive probes correctly, Intel ANS starts the probe retry mechanism. The retry mechanism delays the probe check time until all retries have been sent. The number of retries can be configured using the "Max no. times probe sent" setting found under the team's advanced tab in the device manager. This setting defines the number of times a probe will be sent to an adapter, without a response from the adapter, before the adapter is marked disabled in the team.

Alternatively, the probe burst mechanism allows Intel ANS to send multiple bursts of probe packets. Sending bursts decreases the probability of single probes being lost due to stress on the switch or other factors.

Probe packets used in this method are reserved as "type" 0x886D. Within the 0x886D probe packets are ID fields that indicate what team has sent the probe and specifically what member of the team has sent the probe. This prevents probes from different teams from interacting with one another, and it also tells the receiver of the probe which adapter in the team sent the probe.

Use of probe packets is supported by the AFT, SFT, and ALB/RLB teaming modes. Detection by probe packets is not supported for either of the link aggregation modes (SLA or 802.3ad dynamic).

Activity-Based Tolerance

A second method for determining network connectivity at the server adapter is detection of network activity at the adapter, or traffic detect. In this method, the teamed adapter detects whether it is receiving traffic on the network, such as server broadcasts. While broadcasts come from other nodes in the network, rather than from the server connections, broadcast activity is a valid indication of network connectivity. This method is supported by all teaming modes.

Link-Based Tolerance

Link-based tolerance is the third method for determining whether network connections are in an active state. With the link-state method, the operating system (OS) reports back whether a given network connection has link or not. Since adapters rely on the OS to report link status, some network administrators prefer to use link-based tolerance rather than probe packets in order to cut down on the number of packets on their network. The reasoning is that by using link-based tolerance, instead of using network bandwidth to send and receive probe packets, there will be more bandwidth for real network traffic.

Also, the SLA and 802.3ad dynamic link-aggregation modes do not support probe packets, but do support link-failure detection by link-based tolerance. With these teaming modes, all member ports in the team have the same MAC address, so probe packets cannot be directed to each adapter in the team. As result, some other detection mechanism must be used, such as link-based tolerance.

Whatever detection method is used—whether it is probe packets, activity-based tolerance, or link-based tolerance—a failure-to-detect results in failover of network traffic to one of the available standby adapters in the team. This Adapter Fault Tolerance mechanism is the most basic function in adapter teaming and is included as the basis of the other teaming functions (ALB and SLA) in the teaming hierarchy.

Server Load-Balancing Mechanisms

There are two modes of load balancing supported by Intel ANS. These are Adaptive Load Balancing (ALB) and Receive Load Balancing (RLB), which is actually a subset of ALB. ALB provides load balancing of transmit (outbound) traffic. By default, ALB also includes RLB. Together these two teaming modes permit load balancing in both transmit and receive directions. RLB can be independently disabled. The ALB/RLB teaming modes work with any switch, and AFT is inherent in the ALB/RLB modes.

Adaptive Load Balancing

In Adaptive Load Balancing (ALB) transmit traffic is balanced across the connections in the team to increase transmit (Tx) throughput. In transmit load balancing, a hash table is used to assign flows destined for a particular end client to one of the ports in the team.

The hash algorithm uses a Layer 3 hash index derived from the last octet of the end client's IP address. The transmit traffic from the server is then transmitted through the member in the team corresponding to that index in the hash table (see Table 1). New data flows from the server are assigned to the least loaded member in the team, and this member is placed in the table corresponding to the client hash index. The clients will continue to receive traffic from that particular team member in the server until a pre-configured load balancing interval timer expires. When the interval timer expires, data flows are rebalanced among the network connections in the team.

ALB does not balance traffic on the server receive (Rx) side. Instead, all traffic is received by the server on the primary adapter in the team.

In ALB mode (without Receive Load Balancing enabled), the MAC address of the primary team member is used in the Address Resolution Protocol (ARP) reply packet, and is not load balanced. This ensures that clients learn the MAC address of the primary adapter, and all receive traffic is, hence, directed to the primary in the team. All other packets transmitted from the teamed ports are load balanced and contain the MAC address in the Ethernet header of the team member doing the transmitting .

IP/MAC Address	Adapter 0	Adapter 1	Adapter 2	Adapter 3
00	900			
01		1500		
02				
-				
-			1125	
-				
FF				750

Table 1. Load Balancing Algorithm

In a failover situation, a secondary network connection will receive the MAC address of the primary network connection, the switch MAC address table will be updated, and traffic will be received by the new primary network connection. Therefore, it is important that, if the primary network connection is "Hot Plugged" out of the server, the adapter is not inserted somewhere else in the network unless the server is rebooted. Upon reboot, the team will take on the MAC address of the new primary adapter.

Protocols supported under the ALB scenario are TCP/IP and UDP/IP (both directions, all supported frame types). With NetBEUI and IPX on Windows-based servers, traffic is transmitted and received on the primary network connection only (see Table 2). The reason for Rx on the primary connection only is that NetBEUI is not routable as a protocol and Windows does not route IPX on Windows-based servers. However, on NetWare-based servers, IPX and TCP/IP packets are balanced across network connections in the team.

Adapter Fault Tolerance and Switch Fault Tolerance

OS	Fail-over/Fallback—Intel				Fail-over/Fallback—Multi-vendor			
	IPv.4	IPv.6	IPX	NetBEUI	IPv.4	IPv.6	IPX	NetBEUI
Win 2K	Y	N/S	Y	Y	Y	N/S	Y	Y
Win 2K3	Y	Y	Y	Y	Y	Y	Y	Y
Win XP	Y	Y	Y	Y	Y	Y	Y	Y
NW 6.5	Y	Y	Y	N/S	Y	Y	Y	N/S

Adaptive Load Balancing and Receive Load Balancing

OS	Fail-over/Fallback—Intel				Fail-over/Fallback—Multi-vendor			
	IPv.4	IPv.6	IPX	NetBEUI	IPv.4	IPv.6	IPX	NetBEUI
Win 2K	Y	N/S	Y	Y	Y	N/S	Y	Y
Win 2K3	Y	Y	Y	Y	Y	Y	Y	Y
Win XP	Y	Y	Y	Y	Y	Y	Y	Y
NW 6.5	Y	Y	Y	N/S	Y	Y	Y	N/S

Static Link Aggregation and 802.3ad Dynamic Mode

OS	Fail-over/Fallback—Intel				Fail-over/Fallback—Multi-vendor			
	IPv.4	IPv.6	IPX	NetBEUI	IPv.4	IPv.6	IPX	NetBEUI
Win 2K	Y	N/S	Y	Y	Y	N/S	Y	Y
Win 2K3	Y	Y	Y	Y	Y	Y	Y	Y
Win XP	N/S	N/S	N/S	N/S	N/S	N/S	N/S	N/S
NW 6.5	Y	Y	Y	N/S	Y	Y	Y	N/S

Table 2. Supported Protocols for x86, x64, and Itanium® Platforms

Receive Load Balancing

Receive load balancing (RLB) is a subset of ALB. It allows traffic to flow in both Tx and Rx on all adapters in the team. When creating an RLB team in Windows, this feature is turned on by default. It can be disabled via the Intel® PROSet GUI using the team's Advanced Settings.

In RLB mode, when a client is trying to connect to a team by sending an ARP request message, Intel ANS takes control of the server ARP reply message coming from the TCP stack in response. Intel ANS then copies into the ARP reply the MAC address of one of the ports in the team chosen to service the particular end client, according to the RLB algorithm. When the client gets this reply message, it includes this match between the team IP and given MAC address in its local ARP table. Subsequently, all packets from this end client will be received by the chosen port. In this mode, Intel ANS allocates team members to service end-client connections in a round-robin fashion, as the clients request connections to the server.

In order to achieve a fair distribution of end clients among all enabled members in the team, the RLB client table is refreshed at even intervals (default is five minutes). This is the Receive Balancing Interval, which is a preconfigured setting in the registry. The refresh involves selecting new team members for each client as required. Intel ANS initiates ARP Replies to the affected clients with the new MAC address to connect to, and redistribution of receive traffic is complete when all clients have had their ARP tables updated by Intel ANS.

The OS can send out ARP requests at any time, and these are not under the control of the Intel ANS driver. These are broadcast packets sent out through the primary port. Since the request packet is transmitted with the team's MAC address (the MAC address of the primary port in the team), all end clients that are connected to the team will update their ARP tables by associating the team's IP address with the MAC address of the primary port. When this happens, the receive load of those clients collapses to the primary port.

To restart Rx load balancing, Intel ANS sends a gratuitous ARP to all clients in the receive hash table that were transmitting to non-primary ports, with the MAC address of the respective team members. In addition, the ARP request sent by the OS is saved in the RLB hash table, and when the ARP reply is received from the end client, the client's MAC address is updated in the hash table. This is the same mechanism used to enable RLB when the server initiates the connection.

Since RLB requires that each member in the team receive packets from different clients, all members of the team in RLB mode use their own permanent MAC addresses as the source MAC in the Ethernet header for transmitted packets. Hence, the corresponding switch ports learn the MAC address of each individual member, allowing the switch ports to forward packets from end clients to the appropriate member of the team.

In the Windows Operating System, on failover from the primary to a secondary, the intermediate driver software puts the secondary port into promiscuous mode, and a filter with the primary port's MAC address is applied to it. The primary adapter's virtual MAC address is set to 00:00:00:00:00:00 and the secondary port, upon assuming the primary adapter's MAC address, sends a learning packet of type 0x6D88 to the switch so that the primary adapter's hardware address is learned on the secondary adapter's switch port. This is to protect against receive-packet loss from clients that were connected to the primary adapter. When the subsequent receive balance interval timer expires, the appropriate secondary MAC address will be sent to the team's clients via an ARP reply to ensure that the receive load is once again balanced across all active ports in the team.

Protocols supported under RLB are the same as ALB, as shown in Table 2 on the previous page.

Static Link Aggregation (SLA)

Static Link Aggregation (SLA) is a teaming mode where all of the ports in the team share the same MAC address and are perceived as a single link from the switch's perspective. This is a switch-controlled teaming mode. With Cisco switches, the channeling protocol on the module that the team is plugged into must be set to PAgP, and the channeling mode must also be set to "on". Non-Cisco devices should be configured per the switch vendor's load-sharing instructions. In all cases, SLA is a static mode in the sense that the ports on the switch are active by default and no handshake/negotiation takes place between the switch and the intermediate driver.

There is no designated primary port in the team (all ports are equal), and all ports send and receive simultaneously. Intel ANS does not implement any failover mechanism for receive-side traffic in this mode since there are no MAC address and packet-filter changes. Failover of receive-side traffic is handled entirely by the switch, and Intel ANS implements the transmit load-balancing mechanism in a manner similar to the ALB mode.

When configuring an SLA team, take care to ensure that all the links are of the same speed. Since the switch handles receive-side load balancing, mixed-speed teams may not be effective and may not be supported by some switches.

When removing an existing team member from the team or when adding a new member, care should be taken that the corresponding switch ports are configured to be part of the existing channel or removed from the channel. Additionally, to prevent packet loss, adapters being added or removed from an SLA team should be unplugged from the switch ports prior to changing the team or switch configuration.

802.3ad Dynamic Mode (IEEE 802.3ad)

Dynamic Link Aggregation refers to the dynamic IEEE 802.3AD teaming mode. This is similar to SLA in the sense that all members of this team share the same MAC address and appear as one large link to the switch. This is also a switch-controlled mode; however, in this mode, the corresponding teamed ports on the switch are configured to use the LACP protocol, and are set to mode "ACTIVE". This configuration allows the switch ports to establish dynamic communication with the Intel ANS intermediate driver and allows the user to dynamically add or remove members from the team without concern for packet loss, unless it is the "primary port" as described below.

There is no designated primary in the team (user configured or chosen by Intel ANS). However, the first teamed port on the switch is treated as the initiator by the switch. Hence, removal of this primary member from the team could lead to packet loss. To avoid this scenario, always make sure to pre-configure the switch ports for added or removed members before making any changes to the team itself. Check the documentation for the specific switch to determine how to configure the ports.

When members of mixed speeds are configured in this teaming mode, the switch splits the team into separate channels, one channel each for the aggregate of members of matching speed. The user can control which channel is active on the switch by choosing either the "Maximum Adapters" or the "Maximum Bandwidth" setting from the Advanced Page of the team in PROSet. Failover to the inactive channel, in this case, occurs only when all the links in the active channel malfunction. However, it should be noted that mixed-speed IEEE 802.3ad teams may not be supported by some switches.

Multi-Vendor Teaming

Intel ANS also allows teaming with third-party adapters and LOMs. Some Original Equipment Manufacturers (OEMs) ship several vendors' adapters and networking silicon, and Intel supports these network connections as members of teams. However, from a practical standpoint, not every adapter and piece of LAN silicon on the market can be supported because the testing required for ensuring stable functionality for every product would result in an impossibly long test cycle. As a result, the vendors supported by Intel ANS are limited to third-party network connections that some OEMs currently ship on their servers. Tests with several other LAN adapters have resulted in system instability or caused blue-screens. By testing and supporting a limited set of network connections, Intel can ensure that customers will be covered with supported solutions.

The main restriction of Multi-Vendor Teaming (MVT) with Intel ANS is that VLANS are not supported on teams containing non-Intel adapters.

Network Topologies

There are many ways to configure servers and a network for teaming. For example, the switches used can be either Layer 2 or Layer 3 Routing Switches. Figure 4 also offers some examples for deploying teaming.

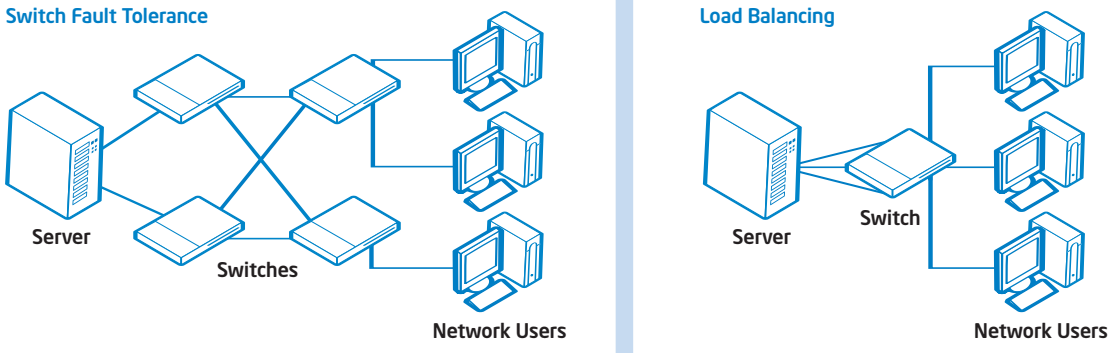

Figure 4. Adapter Teaming Topologies

The left side of Figure 4 shows use of two teamed adapters with multiple cross-connected switches to provide switch fault tolerance. This topology has been discussed previously and requires use of Spanning Tree Protocol to avoid looping problems. The advantages are that the network is both adapter fault tolerant and switch fault tolerant. Additionally, the network users are segmented across two switches, offering increased network capacity over that of a single-adapter, single-switch combination serving one network segment.

The right side of Figure 4 shows multiple teamed adapters connected to a single switch. This configuration allows teaming of up to eight Ethernet ports in any of the hierarchy of teaming modes, from Adapter Fault Tolerance to Static or Dynamic Link Aggregation. The chief advantages of this topology are the inclusion of Adapter Fault Tolerance along with scaling up of bandwidth as more adapters are teamed for link aggregation

Conclusion

Intel has been providing Advanced Network Services features since 1996 and continues to add features for increased server uptime and server scalability. Table 3 summarizes the features now available with Intel ANS software.

Additionally, Intel has made contributions to other teaming software, such as the Open Source Linux Channel Bonding driver. We will continue to enhance Intel ANS with new features in the future. We hope that this document helps explain how Intel ANS software works so that you feel more comfortable deploying these features in your enterprise network. Thank you for choosing Intel PRO Network Connections.

Function	Intel (AFT, SFT, ALB)	Intel (SLA)	Intel (802.3ad)
Number of NICs per team	8 (2 for SFT)	8	8
NIC Fault Tolerance	Yes	Yes	Yes
Switch Fault Tolerance	2	Yes	Yes
Tx Load Balance	Yes	Yes	Yes
Rx Load Balance	Yes	Yes	Yes
Requires compatible switch	Yes	Yes	Yes
Heartbeats to check connectivity	Yes	No	No
NICs with different media	Yes	No	Yes
NICs with different speeds	Yes	No	Yes
Load balances TCP/IP	Yes	Yes	Yes
Load balances other protocols	Yes	Yes	Yes
Same MAC address for all team members	No	Yes	Yes
Same IP address for all team members	Yes	Yes	Yes
Load balancing by IP address	Yes	No	No
Load balancing by MAC address	Yes	Yes	Yes
802.1q tagged VLANs	Yes	Yes	Yes
Untagged VLANs	Yes	Yes	Yes

Table 3. Intel® ANS Feature Comparison

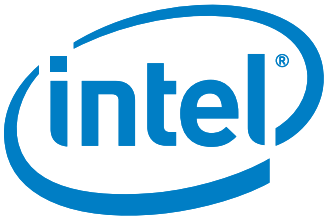

For more information on Intel® ANS software, visit
www.intel.com/support/network/adapter/ans/

For information on the full line of Intel® PRO connectivity products, visit
www.intel.com/network

0606/FM/PMS/PP/1K 312201-001US