COMPAQ
Whitepaper

January 2002
162M-0102A-WWEN

Prepared by Industry Standard Server Group (Austin Development Group)

Compaq Computer Corporation

Contents

Compaq Network Adapter Teaming Technology

Abstract This document addresses the teaming technology behind Network Fault Tolerance (NFT), Transmit Load Balancing (TLB) and Switch-assisted Load Balancing (SLB), including failure recovery methods, load balancing logic, and network scenario considerations.

NOTICE

Compaq Network Adapter Teaming Technology

Whitepaper prepared by Industry Standard Server Group (Austin Development Group)

First Edition (January 2002)
Document Number 162M-0102A-WWEN

Introduction

This document provides detailed information about the design, implementation, and configuration of Compaq's network adapter teaming, which includes network fault tolerant and load balancing technology. Although this document specifically discusses the teaming of Compaq network adapters under Microsoft Windows 2000, many of the concepts of Compaq Network Adapter Teaming are applicable to other operating systems.

The design goal of Compaq's Network Adapter Teaming is to provide fault tolerance and load balancing across a team of two or more network adapters. The term "team" refers to the concept of multiple network adapters working together as a single network adapter, commonly referred to as a Virtual Network Adapter.

The purpose of this document is to assist networking specialists, systems engineers, and IM professionals in the design and troubleshooting of environments incorporating this technology in Compaq servers. This whitepaper assumes that the reader is familiar with the basics of IP, the OSI model, the use of network drives, and the fundamentals of network switching. Additionally, the reader should be familiar with the terms found in the section titled, "Definitions and Acronyms" in this whitepaper.

Note: Information in this document was derived from network driver version NCDE6.21 (August 16, 2001). Since the white paper is about technology, most of the information is generally applicable to future releases; however, specific features may differ slightly between revision levels.

Overview of Network Addressing

Understanding the concepts of network addressing is key to understanding how Compaq's Network Adapter Teaming works. This section provides a brief overview of network addressing as a baseline for explaining how Compaq's Network Adapter Teaming can create one logical network adapter from a team of two or more adapters.

Layer 2 vs. Layer 3 Addressing

Devices on a computer network use unique addresses, much like telephone numbers, to communicate with each other. Each device, depending on its function, will use one or more of these unique addresses. The addresses correspond to one or more layers of the OSI model. Most often, network devices use an address at Layer 2, the Data Link Layer, called a MAC address, and an address at Layer 3, the Network Layer, called a protocol address (e.g., IP, IPX, AppleTalk). One could say that a MAC address is one that is assigned to the hardware, whereas a protocol address is one that is assigned to the software.

MAC addresses are in the format of 00-00-00-00-00-00 (hexadecimal), IP addresses in the format of 0.0.0.0 (dotted decimal), and IPX addresses in the format of 000000.000000000000 (hexadecimal). Since multiple protocols can reside on the same network device, it is not uncommon for a single network device to use one MAC address and one or more protocol addresses.

Ethernet devices communicate directly using the MAC address, not the protocol address. For instance, when a PING is initiated for the address 1.1.1.1, the network device must find a corresponding MAC address for the IP address of 1.1.1.1. A frame is then built using the destination MAC address and sent to the destination computer. The frame carries the sender's protocol address in its payload, which is how the destination network device knows to which device to respond. This means that protocol addresses must be resolved to MAC addresses. For IP, this is done using ARP. (refer to section titled, " Scenarios of Network Addressing and Communication") For IPX, the second part of the IPX address is the same as the hardware address, so no special mechanism is needed.

Addresses: Unicast vs. Broadcast vs. Multicast

There are three types of Layer 2 and Layer 3 addresses: unicast, broadcast, and multicast. A unicast address is one that corresponds to a single network device, either a single MAC address or a single IP address. When a station transmits a frame to a unicast address, the transmitting device intends for only a single network device to receive the frame. When a station transmits a frame to a broadcast MAC address or IP address, the station intends for all devices on a particular network to receive the frame. When a station transmits a frame to a multicast MAC or IP address, the station intends for a predefined group of network devices to receive the frame. A group, as used here, can be defined as more than one network device but less than all the network devices on a particular network.

A Multicast address is used in Compaq Network Adapter Teaming for the purpose of transmitting and receiving heartbeat frames (refer to section titled, " Heartbeats").

Compaq Network Adapter Teaming and Layer 2/Layer 3 addresses

One of the most important concepts to understand when implementing Compaq Network Adapter Teaming is that of Layer 2 and Layer 3 addresses and the way they are handled. When teaming network adapters together, they function as a single virtual network adapter. Other network devices communicating with a Compaq Network Adapter Team cannot distinguish that they are communicating with more than one network adapter. Compaq Network Adapter Teaming must maintain strict IEEE standards compliance in its use of Layer 2 and Layer 3 addresses.

In order for a Compaq Network Adapter Team to appear as a single virtual network adapter, it is necessary for all networking devices to refer to the team by a single Layer 2 address and a single Layer 3 address. In other words, when a device is communicating with a team, regardless of the number of network adapters that make up the team, the network device only "sees" one MAC address and one protocol address (e.g., IP, IPX). When communicating using IP, this means that a networking device will have only one entry in its ARP cache for a Compaq Network Adapter Team regardless of the number of network adapters that make up the team.

When a Compaq Network Adapter Team initializes at server boot time, the Teaming driver for each team (up to 16 teams of eight network adapters each may exist in a single server) "reads" the BIA for each network adapter assigned to that particular team. Essentially, the MAC addresses are decoupled from the network adapters and pooled together for use by the Teaming driver. The Teaming driver picks one MAC address as the Team's MAC address and assigns it to the Primary Adapter, unless the user has manually set the MAC address (Locally Administered Address) via the configuration GUI (Compaq Teaming and Configuration Utility). In addition, all ARP Replies from the server for this particular Compaq Network Adapter Team provide this same MAC address as the team's MAC address. This address does not change unless the team is reconfigured and rebooted. The Teaming driver assigns the remaining MAC addresses to the Secondary Adapters.

When a failover event occurs, the MAC addresses of the current Primary Adapter and one of the Secondary Adapters are swapped. The former Secondary Adapter becomes the new Primary Adapter and the former Primary Adapter becomes a Secondary Adapter. By swapping the MAC addresses in this manner, the Compaq Network Adapter Team is always known by one MAC address and one Protocol address. It is unnecessary for Protocol addresses to swap during a failover event, because the protocol address is directly assigned to the Intermediate (Teaming) driver, and not to the Miniport driver.

When transmitting frames, the current Primary Adapter always transmits using the Team's MAC address as the Layer 2 address and the Team's Protocol address as the Layer 3 address. Secondary Adapters always transmit using the MAC address assigned to them by the Teaming driver and using the Team's protocol address as the Layer 3 address. For NFT and TLB, the MAC address used when transmitting is always different from the Primary Adapter's MAC address and is always unique from any other Secondary Adapter, for IEEE standards compliance. For SLB, the additional switch intelligence allows all teamed adapters to transmit using the same MAC address, the Team's MAC address.

Network device communicating with a Compaq Network Adapter Team may receive frames from more than one network adapter in the same team. When this happens, the network device does not know that more than one Layer 2 address is being used. The important issue is that all frames originating from the same Compaq Network Adapter Team use the same Protocol address. The network device does not know that multiple MAC addresses are coming from the Team because MAC headers are stripped off before the frames are processed up the stack by the network device's operating system. By the time the operating system receives the frames, they all appear as though they came from the same network adapter. In addition, ARP cache entries are not made by learning the Layer 2 addresses from received frames. ARP cache entries are ONLY made from ARP Replies or by static entries by hand. Since the Team always sends ARP Replies using the same MAC address, the Team is only known by one MAC address to all network entities.

Scenarios of Network Addressing and Communication

As discussed earlier, protocol addresses (e.g., IP, IPX) must be resolved to hardware addresses (MAC) for network devices to communicate. What follows are two simple scenarios with one network device (named Red) PINGing another network device (named Blue). The first scenario cites one device PINGing another on the same Layer 2 network. The second scenario cites one device PINGing another on a different Layer 2 network, which requires the use of a router to effect communication.

These scenarios provide a baseline of typical network addressing and communication using IP. This baseline will be referred to later in this document to differentiate how Compaq Network Adapter Teaming functions in these same scenarios. By understanding the differences in simple examples such as these (without teaming technology involved), implementers will have a better understanding of how this technology may work in their environment.

Scenario 1: Device PINGs another on the same Layer 2 network. *(See Figure 1.)*

1. Red transmits a broadcast ARP Request asking for Blue's MAC address.

   A user on Red issues the command "ping 1.1.1.2" to initiate a PING to Blue. The number 1.1.1.2 refers to Blue's IP address, or protocol address. First, Red determines whether or not Blue is on the same Layer 2 network by running an algorithm using its own IP address of 1.1.1.1, its own subnet mask (not shown), and Blue's IP address of 1.1.1.2. If Blue is on a different Layer 2 network, then Red will need to use its Gateway, or router, to get to Blue.

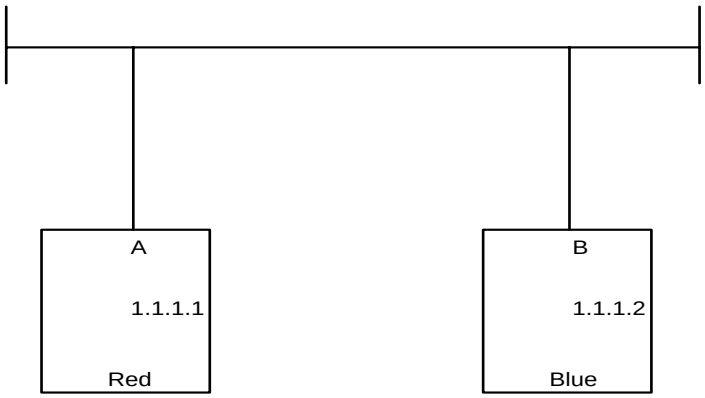

Figure 1. One device PINGs another on the same Layer 2 network

   Once Red has determined that Blue is on the same Layer 2 network, Red must find out what Blue's MAC address is. First, Red checks its own ARP cache for a MAC address entry matching the IP address of 1.1.1.2. ARP is used to map protocol addresses to hardware addresses. If Red does not have a static entry or an entry cached from a previous conversation with Blue, then it must broadcast an ARP Request frame containing the IP address of Blue on the network asking Blue to respond and provide its MAC address. Red must broadcast this ARP Request because without knowing Blue's unique MAC address, it has no way of sending a frame directly (unicast) to Blue.

2. Blue transmits a unicast ARP Reply to Red, providing its MAC address (B).

   Blue sees the ARP Request containing its own IP address and responds with a unicast ARP Reply directly to Red. Blue also notes Red's MAC address (A) and IP address of 1.1.1.1, and enters them into its ARP cache. Red receives the ARP Reply and enters Blue's MAC address (B) IP address of 1.1.1.2 into its own ARP cache.

3. Red transmits a unicast PING Request to Blue using Blue's destination MAC address of 1.1.1.2 (B).

Red can now create a PING Request frame using Blue's MAC address (B). Red sends the PING Request to Blue using Blue's destination MAC address of 1.1.1.2 (B). Blue receives the PING Request frame and notices that a station with an IP address of 1.1.1.1 is requesting that it respond.

Note: The following step may not occur if Blue's ARP table still contains an entry for Red as a result of steps 1 and 2.

4. Blue transmits a broadcast ARP Request asking for Red's MAC address

   Blue checks its ARP cache for a MAC address entry that matches 1.1.1.1. If Blue does not find one (i.e., ARP cache timed out since last communication with Red), then Blue broadcasts an ARP Request asking for Red's MAC address.

5. Red transmits a unicast ARP Reply to Blue providing its MAC address (A).

   Note: This step may not occur if step 4 does not take place.

   Red sees the ARP Request and transmits a unicast ARP Reply directly to Blue providing its MAC address (A). Blue receives the ARP Reply and puts Red's MAC address (A) and IP address (1.1.1.1) in its ARP cache.

6. Blue transmits a unicast PING Reply to Red using Red's destination MAC address of 1.1.1.1 (A).

   Blue transmits a unicast PING Reply to Red using Red's MAC address (A) and the user sees the PING REPLY message printed on the screen. This completes the entire conversation.

Scenario 2: Device PINGs another on a different Layer 2 network. *(See Figure 2)*

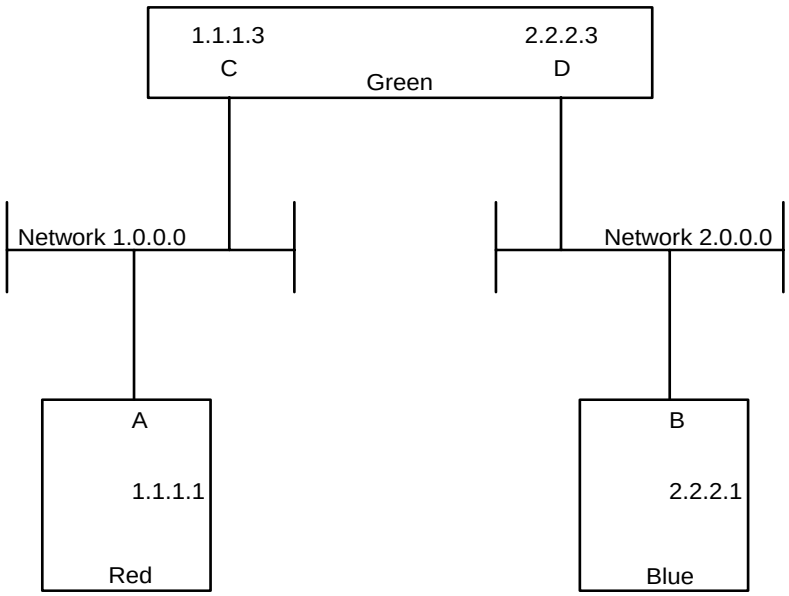

Figure 2. One device PINGs another on a different Layer 2 network

1. Red transmits a broadcast ARP Request on Network 1.0.0.0 asking for Green's MAC address.

A user on Red issues the command "ping 2.2.2.1" to initiate a PING to Blue. The number 2.2.2.1 refers to Blue's IP address, or protocol address. First, Red determines whether or not Blue is on the same Layer 2 network by running an algorithm using its own IP address of 1.1.1.1, its own subnet mask (not shown), and Blue's IP address of 2.2.2.1. If Blue is on a different Layer 2 network, then Red will need to use its Gateway (Green), or router, to get to Blue.

Once Red has determined that Blue is on a different Layer 2 network, Red must use Green as a Gateway to get to Blue. Red communicates directly with Green at Layer 2 but communicates directly with Blue at Layer 3. This means that Red must transmit a unicast frame with the Layer 2 address (MAC) of Green, but the same frame will have Blue's Layer 3 address in it. When Green receives the frame, it sees the Layer 3 data destined for Blue and forwards the frame onto Blue via the Green's interface that is attached to Blue's Layer 2 network. This means that Red must find out what Green's MAC address is. First, Red checks its own ARP cache for an entry that matches 1.1.1.3 (Red's Gateway). If Red doesn't have an entry cached, then it must broadcast an ARP Request frame on the network asking Green to respond and provide its MAC address.

2. Green transmits a unicast ARP Reply to Red providing its MAC address (C).

   Green sees the ARP Request and responds with a unicast ARP Reply to Red. Also, Green enters Red's MAC address and IP address into its ARP cache. Red receives Green's reply and enters the MAC address and the IP address of Green into its ARP cache.

3. Red transmits a PING Request to Blue (2.2.2.1) using the destination MAC address of Green's 1.1.1.3 interface (C), since Green is Red's Gateway to Blue.

   Red can now create a PING Request frame using Green's MAC address and Blue's IP address. Red sends the PING Request. Green receives the PING Request and determines that the frame is meant for Blue because of the Layer 3 address (IP).

4. Green transmits a broadcast ARP Request on Network 2.0.0.0 asking for Blue's MAC address.

   Green looks in its ARP cache for a MAC address for Blue. If one is not found, Green broadcasts an ARP Request frame on Blue's Layer 2 network asked for Blue's MAC address.

5. Blue transmits a unicast ARP Reply to Green providing its MAC address (B).

   Blue sees the ARP Request frame and responds with a unicast ARP Reply frame to Green. Also, Blue enters Green's MAC address and IP address into its ARP cache. Green receives the ARP Reply from Blue and enters Blue's MAC address and IP address into its ARP cache.

6. Green forwards Red's PING Request to Blue using Blue's destination MAC address (B).

   Green now transmits Red's Ping Request frame onto Blue's network using Blue's MAC address and Blue's IP address as the destination MAC and destination IP address. The source MAC address is Green's MAC address and the source IP address is Red's IP address. Blue receives the frame and notices that a station with an IP address of 1.1.1.1 is asking for it to respond to a PING. Before Blue can respond with a PING Reply, it must determine whether or not 1.1.1.1 is on the same layer 2 network. Blue runs an algorithm (details are beyond the scope of this document) using its own IP address (2.2.2.1), its own subnet mask (not shown) and the IP address of Red (1.1.1.1). Blue then determines that Red is on a different network. Because of this, Blue must use its gateway (Green) to get the PING Reply back to Red.

7. Blue transmits a broadcast ARP Request on Network 2.0.0.0 asking for Green's MAC address.

Note: This step may not occur if Blue's ARP table still contains an entry for Red resulting from steps 4 and 5.

Blue checks its ARP cache for the MAC address that corresponds to the IP address of 2.2.2.3 (Blue's Gateway). If an entry isn't found, Blue must broadcast an ARP Request asking for Green's MAC address.

8. Green transmits a broadcast ARP Reply to Blue providing its MAC address (D).

   Note: This step may not occur if Blue's ARP table still contains an entry for Red resulting from steps 4 and 5.

   Green sees the ARP Request and responds with a unicast ARP Reply directly to Blue. Also, Green enters Blue's MAC address and IP address into its ARP cache. Blue receives the ARP Reply and puts Green's MAC address and IP address in its ARP cache. Blue now has all the information it needs to send a PING Reply to Red.

9. Blue transmits a unicast PING Reply to Red (1.1.1.1) using the MAC address of Green's 2.2.2.3 interface (D).

   Blue transmits a unicast PING Reply to Red by way of Green by using Green's MAC address as the destination MAC address, Red's IP address as the destination IP address, Blue's MAC address as the source MAC address and Blue's IP address as the source IP address. Green receives the PING Reply and determines that the frame is meant for Red because of the Layer 3 address (IP).

10. Green transmits a broadcast ARP Request on Network 1.0.0.0 asking for Red's MAC address.

    Note: This step may not occur if Green's ARP table still contains an entry for Red resulting from steps 1 and 2.

    Green looks in its ARP cache for a MAC address for Red. If one is not found, Green broadcasts an ARP Request frame on Red's Layer 2 network asking for Red's MAC address.

11. Red transmits a unicast ARP Reply to Green providing its MAC address (A).

    Note: This step may not occur if step 10 does not take place.

    Red sees the ARP Request frame and responds with a unicast ARP Reply frame to Green. Also, Red enters Green's MAC address and IP address into its ARP cache. Green receives the ARP Reply from Red and enters Red's MAC address and IP address into its ARP cache.

12. Green forwards Blue's PING Reply to Red using the destination MAC address of Red (A).

    Green transmits Blue's Ping Reply frame onto Red's network using Red's MAC address and Red's IP address as the destination MAC and destination IP address. The source MAC address is Green's MAC address and the source IP address is Blue's IP address. The user sees the PING REPLY message printed on the screen. This completes the entire conversation.

Teaming Mechanisms

Architecture of Compaq Network Adapter Teaming

Within an operating system (OS), a hierarchy of layers work together to enable one OS to communicate with another. Each of these layers performs a separate function and passes information between the layers above and below it. Within Windows 2000, there are four layers that are important when discussing Compaq Network Adapter Teaming: the Miniport layer, Intermediate layer, NDIS layer, and Protocol layer.

- **Miniport Layer**

  The network adapter driver resides at the Miniport Layer. Typically, this driver is written by the vendor of the network adapter. Compaq network adapters drivers (e.g., N100NT5.SYS) are considered Miniport drivers.

- **Intermediate Layer**

  The Intermediate layer driver provides a network function, but is not considered a Miniport since it does not directly control a piece of hardware. The Intermediate layer driver performs a function that is between the Miniport layer and NDIS. The networking function that is performed by the Intermediate layer is beyond the ability of a Miniport layer driver. In this case, Compaq Network Adapter Teaming is considered an Intermediate driver (i.e. CPQTEAM.SYS). It performs the function of making several Miniport drivers (network adapter drivers) seamlessly work as a single network adapter that interfaces with NDIS.

- **NDIS**

  NDIS, Microsoft's Network Driver Interface Specification, handles communications between the underlying layers, either Miniports or Intermediates, and the Protocol layer.

- **Protocol Layer**

  The Protocol layer is where common protocols such as IP, IPX, and AppleTalk, interface with NDIS. ((refer to section of Figure 5, labeled, "Windows 2000 Server.")

In the absence of an Intermediate driver, a protocol address is usually assigned to each individual Miniport driver. However, when utilizing Compaq's Network Adapter Teaming, the protocol address is assigned to a single Compaq Network Adapter Teaming instance that represents the underlying Miniports. If more than one Compaq Network Adapter Team exists in a single server, there will be more than one instance of the Compaq Network Adapter Team and an individual protocol address will be assigned to each instance.

Teaming Software Components

Compaq Network Adapter Teaming consists of three components: the Miniport Driver, Intermediate Driver, and configuration GUI.

- **Miniport Driver**

When running a Windows OS, the Miniport driver used with the Compaq network adapter will be either N100NT5.SYS or N1000NT5.SYS depending on the adapter in use. N100NT5.SYS is for all Compaq NC series 10/100 Mbps network adapters and N1000NT5.SYS is for all Compaq NC series 1000 Mbps network adapters.

- **Intermediate Driver**

The Intermediate driver is CPQTEAM.SYS and is used for all teaming functions involving Compaq NC series adapters.

- **Configuration GUI**

The configuration GUI is called the Compaq Teaming and Configuration Utility and the file name is CPQTEAM.EXE. The configuration GUI is launchable from Control Panel or from the Tray icon.

These three components are designed to work as a unit, when one is upgraded; it is advisable to upgrade all components to the current version. For driver updates to Compaq Network Adapter Teaming, please visit Compaq's Network Adapter Driver site at:

http://www.compaq.com/support/files/networking/nics/index.html.

Types of Compaq Network Adapter Teams

There are three teaming modes for Compaq network adapters: Network Fault Tolerance (NFT), Transmit Load Balancing (TLB), and Switch-assisted Load Balancing (SLB). Respectively, each mode gains in features and incorporates most features from the previous teaming mode (see Figure 3). In other words, NFT is the simplest teaming mode, supporting only network adapter fault tolerance. TLB supports adapter fault tolerance plus load balancing of traffic being transmitted from the server. SLB supports adapter fault tolerance, load balancing of traffic being transmitted from the server, plus load balancing of traffic being received by the server.

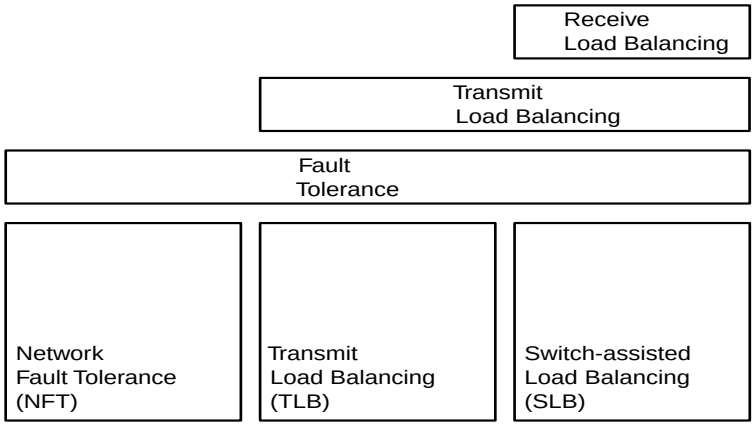

Figure 3. Teaming Types and Teaming Functionality

Network Fault Tolerance (NFT)

Network Fault Tolerance is the most basic form of Compaq Network Adapter Teaming. In this mode, from two to eight adapters may be teamed together as a single virtual network adapter. However, only one network adapter, referred to as the Primary Adapter, is used for both transmit and receive communication with the server. The remaining adapters are considered stand-by adapters, referred to as Secondary adapters, and remain idle until the Primary adapter fails (except for transmitting/receiving heartbeats.)

A Primary Adapter in an NFT team is responsible for all frames transmitted from and received by the server. Secondary Adapters are idle, except for transmitting and receiving heartbeat frames if heartbeats are enabled.

Network Addressing and Communication using NFT

This section builds on the concepts reviewed previously in the section titled, "Scenarios of Network Addressing and Communication", and describes how NFT functions from the network addressing and communication perspective.

Scenario 1:

Utilizing a network diagram similar to Figure 1, Blue has been modified to be a server utilizing a Compaq Network Adapter Team in NFT mode with two network adapters in a team (see Figure XX). The two network adapters have Layer 2 addresses of MAC B and MAC E, respectively, and are known by a single Layer 3 address of 1.1.1.2. Network adapter B has been designated as the Primary Adapter in this NFT team.

1. Red transmits a broadcast ARP Request asking for Blue's MAC address.

   A user on Red issues the command "ping 1.1.1.2" to initiate a PING to Blue. First, Red determines whether or not Blue is on the same Layer 2 network.

   Once Red has determined that Blue is on the same Layer 2 network, Red must find out what Blue's MAC address is. First, Red checks its own ARP cache for a MAC address entry matching the IP address of 1.1.1.2. If Red does not have a static entry or an entry cached from a previous conversation with Blue, then it must broadcast an ARP Request frame on the network asking Blue to respond and provide its MAC address. Red must broadcast this ARP request because without knowing Blue's unique MAC address, it has no way of sending a frame directly (unicast) to Blue.

2. Blue transmits a unicast ARP Reply to Red, providing its MAC address.

   Blue sees the ARP Request (the frame is received on both the Primary and Secondary Adapters in the team) because the frame is broadcasted on the network. However, the team discards all non-heartbeat frames incoming on Secondary Adapters, and responds with a unicast ARP Reply to Red. The ARP Reply is transmitted by the Primary Adapter (B). In Blue's ARP Reply, Blue provides the MAC address of its Teaming driver, which is the same as the current Primary Adapter's MAC address (B) (see Section 0 Compaq Network Adapter Teaming and Layer 2/Layer 3 addresses). Blue also takes note of Red's MAC address (A) and IP address (1.1.1.1) and enters them into its ARP cache. Red receives the reply and enters the MAC address (B) and the IP address of Blue (1.1.1.2) into its own ARP cache.

3. Red transmits a unicast PING Request to Blue using Blue's destination MAC address

Red can now create a PING Request frame using Blue's MAC address (B). Red sends the PING Request to Blue. Blue receives the frame on its Primary Adapter (B) and notices that a station with an IP address of 1.1.1.1 is asking for it to respond.

4. Blue transmits a broadcast ARP Request asking for Red's MAC address.

   Note: The following step may not occur if Blue's ARP table still contains an entry for Red as a result of steps 1 and 2.

   Blue checks its ARP cache for a MAC address entry that matches 1.1.1.1. If Blue does not find one, then Blue broadcasts an ARP Request asking for Red's MAC address.

5. Red transmits a unicast ARP Reply to Blue providing its MAC address.

   Note: The following step may not occur if step 4 does not take place.

   Red sees the ARP Request and transmits a unicast ARP Reply directly to Blue providing its MAC address (A). Blue receives the ARP Reply and puts Red's MAC address (A) and IP address (1.1.1.1) in its ARP cache.

6. Blue transmits a unicast PING Reply to Red using Red's destination MAC address.

   Blue then transmits a unicast PING Reply to Red using Red's MAC address (A) and the user sees the PING Reply message printed on the screen. This completes the entire conversation.

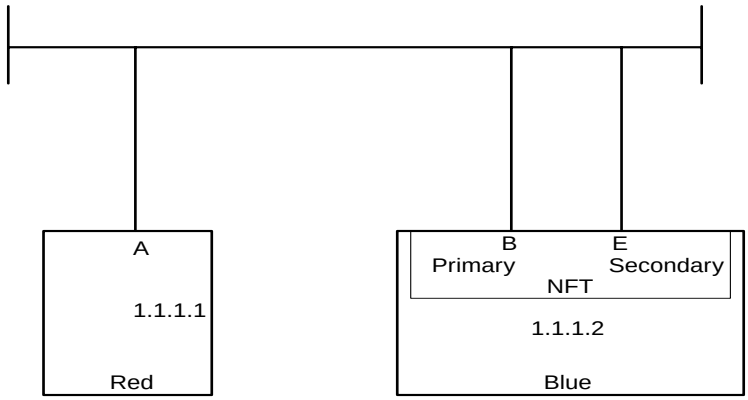

Figure 4. NFT Team PINGs Another Device on the Same Layer 2 Network

NFT Applications

NFT is deployed in environments that only require fault tolerance and do not require transmit or receive throughput greater than the capacity of the Primary Adapter (e.g., a server that requires fault tolerance in

case of a network adapter malfunction, which does not have a demand for receiving or transmitting a large amount of traffic at any given time.)

Recommended Configurations for an NFT Environment

The recommended configuration for an NFT environment is to have all members of the same NFT team attached to the same switch or hub. If switch redundancy is required (team members are attached to two different switches), then it is recommended that the switches be deployed with redundant links between them and Spanning Tree enabled on the ports that connect the switches.

Additionally, Compaq recommends that:

- Heartbeats be enabled (default) and the NFT team MAC address not be manually set to a locally administered address (LAA). A user should not implement LAAs of network adapters that are members of a team, otherwise Teaming may not function correctly.
- Spanning Tree be disabled on all switch ports to which a Compaq Network Adapter Team is attached. Cisco switches have a feature called Port Fast that is used to disable Spanning Tree on a port-by-port basis. If the Compaq-recommended configuration is followed (team members are attached to the same switch), disable Spanning Tree on all switch ports to which the teams are attached.

Transmit Load Balancing (TLB)

Transmit Load Balancing mode, also known as Adaptive Load Balancing (ALB), incorporates all the features of NFT, plus Transmit Load Balancing. In this mode, two to eight adapters may be teamed together as a single virtual network adapter. The load-balancing algorithm used in TLB allows the server to load balance all traffic transmitted from the server. However, traffic received by the server is not load balanced, meaning the Primary Adapter is responsible for receiving all traffic destined for the server (see Figure). In addition, only IP traffic is load balanced. As with NFT, there are two types of team members, Primary and Secondary Adapters. The Primary Adapter transmits and receives frames and the Secondary Adapters only transmit frames.

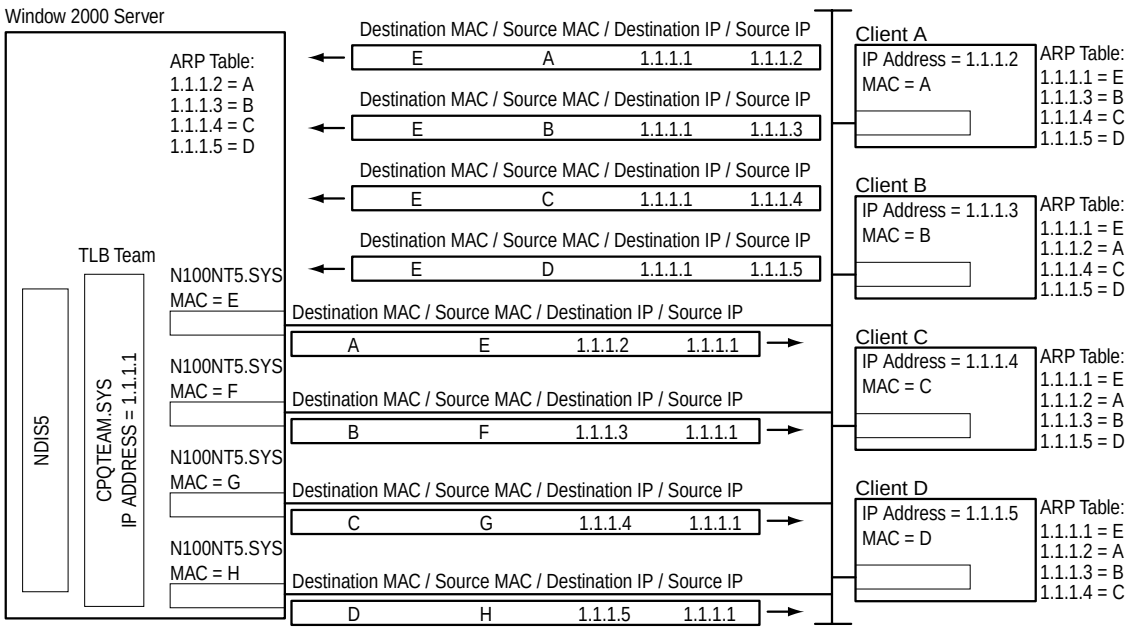

Figure 5. Overview of TLB Communication

Network Addressing and Communication using TLB

Building on the concepts previously reviewed in the section titled, “Compaq Network Adapter Teaming and Layer 2/Layer 3 addresses

One of the most important concepts to understand when implementing Compaq Network Adapter Teaming is that of Layer 2 and Layer 3 addresses and the way they are handled. When teaming network adapters together, they function as a single virtual network adapter. Other network devices communicating with a Compaq Network Adapter Team cannot distinguish that they are communicating with more than one network adapter. Compaq Network Adapter Teaming must maintain strict IEEE standards compliance in its use of Layer 2 and Layer 3 addresses.

In order for a Compaq Network Adapter Team to appear as a single virtual network adapter, it is necessary for all networking devices to refer to the team by a single Layer 2 address and a single Layer 3 address. In other words, when a device is communicating with a team, regardless of the number of network adapters that make up the team, the network device only “sees” one MAC address and one protocol address (e.g., IP, IPX). When communicating using IP, this means that a networking device will have only one entry in its ARP cache for a Compaq Network Adapter Team regardless of the number of network adapters that make up the team.

When a Compaq Network Adapter Team initializes at server boot time, the Teaming driver for each team (up to 16 teams of eight network adapters each may exist in a single server) “reads” the BIA for each network adapter assigned to that particular team. Essentially, the MAC addresses are decoupled from the network adapters and pooled together for use by the Teaming driver. The Teaming driver picks one MAC address as the Team’s MAC address and assigns it to the Primary Adapter, unless the user has manually set the MAC address (Locally Administered Address) via the configuration GUI (Compaq Teaming and Configuration Utility). In addition, all ARP Replies from the server for this particular Compaq Network Adapter Team provide this same MAC address as the team’s MAC address. This address does not change unless the team is reconfigured and rebooted. The Teaming driver assigns the remaining MAC addresses to the Secondary Adapters.

When a failover event occurs, the MAC addresses of the current Primary Adapter and one of the Secondary Adapters are swapped. The former Secondary Adapter becomes the new Primary Adapter and the former Primary Adapter becomes a Secondary Adapter. By swapping the MAC addresses in this manner, the Compaq Network Adapter Team is always known by one MAC address and one Protocol address. It is unnecessary for Protocol addresses to swap during a failover event, because the protocol address is directly assigned to the Intermediate (Teaming) driver, and not to the Miniport driver.

When transmitting frames, the current Primary Adapter always transmits using the Team's MAC address as the Layer 2 address and the Team's Protocol address as the Layer 3 address. Secondary Adapters always transmit using the MAC address assigned to them by the Teaming driver and using the Team's protocol address as the Layer 3 address. For NFT and TLB, the MAC address used when transmitting is always different from the Primary Adapter's MAC address and is always unique from any other Secondary Adapter, for IEEE standards compliance. For SLB, the additional switch intelligence allows all teamed adapters to transmit using the same MAC address, the Team's MAC address.

Network device communicating with a Compaq Network Adapter Team may receive frames from more than one network adapter in the same team. When this happens, the network device does not know that more than one Layer 2 address is being used. The important issue is that all frames originating from the same Compaq Network Adapter Team use the same Protocol address. The network device does not know that multiple MAC addresses coming from the Team because MAC headers are stripped off before the frames are processed up the stack by the network device's operating system. By the time the operating system receives the frames, they all appear as though they came from the same network adapter. In addition, ARP cache entries are not made by learning the Layer 2 addresses from received frames. ARP cache entries are ONLY made from ARP Replies or by static entries by hand. Since the Team always sends ARP Replies using the same MAC address, the Team is only known by one MAC address to all network entities.

Taking the concepts reviewed in the section titled, "Scenarios of Network Addressing and Communication", this section describes how TLB functions from the network addressing and communication perspective.

Scenario 1:

Utilizing a network diagram similar to Figure 1., Blue has been modified to be a server utilizing a Compaq Network Adapter Team in TLB mode with two network adapters in a team (see Figure 6). The two network adapters have Layer 2 addresses of MAC B and MAC E, respectively, and are known by a single Layer 3 address of 1.1.1.2. Network adapter B has been designated as the Primary Adapter in this NFT team.

1. Red transmits a broadcast ARP Request asking for Blue's MAC address.

   A user on Red issues the command "ping 1.1.1.2" to initiate a PING to Blue. First, Red determines whether or not Blue is on the same Layer 2 network.

   Once Red has determined that Blue is on the same Layer 2 network, Red must find out what Blue's MAC address is. First, Red checks its own ARP cache for a MAC address entry matching the IP address of 1.1.1.2. If Red does not have a static entry or an entry cached from a previous conversation with Blue, then it must broadcast an ARP Request frame on the network asking Blue to respond and provide its MAC address. Red must broadcast this ARP request because without knowing Blue's unique MAC address, it has no way of sending a frame directly (unicast) to Blue.

2. Blue transmits a unicast ARP Reply to Red, providing its MAC address.

   Blue sees the ARP Request (the frame is received on both adapters the Primary and Secondary Adapters in the team because the frame is broadcasted on the network. However, the team discards all non-heartbeat frames incoming on Secondary Adapters), and responds with a unicast ARP

Reply to Red. The ARP Reply is transmitted by the Primary Adapter (B) because all non-IP frames are always transmitted by the current Primary Adapter (ARP has an EtherType of 0x0806 and IP has an EtherType of 0x0800).

In Blue's ARP Reply, Blue provides the MAC address of its Teaming driver, which is the same as the current Primary Adapter's MAC address (B) (see Section 0 Compaq Network Adapter Teaming and Layer 2/Layer 3 addresses). Blue also takes note of Red's MAC address (A) and IP address (1.1.1.1) and enters them into its ARP cache. Red receives the reply and enters the MAC address (B) and the IP address of Blue (1.1.1.2) into its own ARP cache.

3. Red transmits a unicast PING Request to Blue using Blue's destination MAC address

   Red can now create a PING Request frame using Blue's MAC address (B). Red sends the PING Request to Blue. Blue receives the frame on its Primary Adapter (B) and notices that a station with an IP address of 1.1.1.1 is asking for it to respond.

4. Blue transmits a broadcast ARP Request asking for Red's MAC address.

   Note: The following step may not occur if Blue's ARP table still contains an entry for Red as a result of steps 1 and 2.

   Blue checks its ARP cache for a MAC address entry that matches 1.1.1.1. If Blue does not find one, then Blue broadcasts an ARP Request asking for Red's MAC address.

5. Red transmits a unicast ARP Reply to Blue providing its MAC address.

   Note: The following step may not occur if step 4 does not take place.

   Red sees the ARP Request and transmits a unicast ARP Reply directly to Blue providing its MAC address (A). Blue receives the ARP Reply and puts Red's MAC address (A) and IP address (1.1.1.1) in its ARP cache.

6. Blue transmits a unicast PING Reply to Red using Red's destination MAC address.

   The final step in the conversation is for Blue to transmit a PING Reply to Red. However, since Blue's Team is running TLB, it must make a load balancing decision before transmitting the PING Reply. The load balancing decision is made by using either Red's MAC address or Red's IP address. Once Blue decides which network adapter to use, it transmits a unicast PING Reply to Red using Red's MAC address (A).

   If Blue chooses the Primary Adapter, Red will receive a PING Reply from Blue with a source MAC address of "B", destination MAC address of "A", a source IP address of "1.1.1.2" and a destination IP address of "1.1.1.1". However, if Blue choose the Secondary, Red will receive a PING Reply from Blue with a source MAC address of "E", destination MAC address of "A", a source IP address of "1.1.1.2" and a destination IP address of "1.1.1.1". Either way, Red only distinguishes that it received a PING Reply from the Layer 3 address, "1.1.1.2" (see section titled, "TLB Transmit Balancing Algorithm", for a complete discussion). The user sees the PING Reply message printed on the screen. This completes the entire conversation.

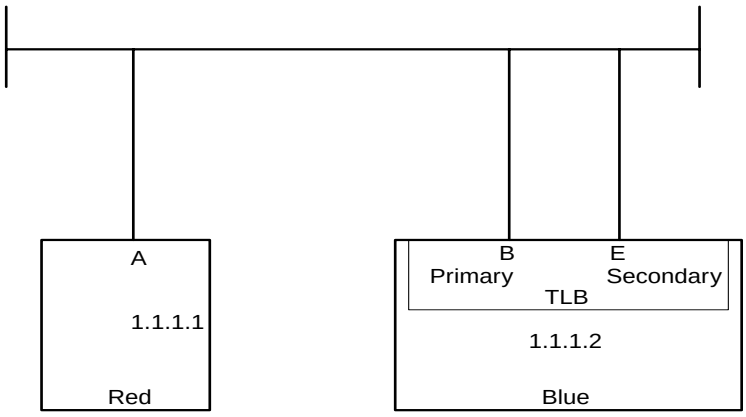

Figure 6. TLB Team PINGs Another Device on the Same Layer 2 Network

TLB Transmit Balancing Algorithm

A Team in TLB mode attempts to load balance transmitted frames. In order to avoid frames being transmitted out of order when communicating with a single network device, the load balancing algorithm assigns conversations to a particular adapter. In other words, load balancing is performed on a conversation-by-conversation basis rather than on a frame-by-frame basis. To accomplish this, when making a decision about which teamed adapter will transmit the frame, the algorithm uses the destination MAC address or the destination IP address of the frame to be transmitted.

It is very important to understand the differences the algorithm uses when deploying Compaq Network Adapter Teaming in an environment that requires load balancing of routed Layer 3 traffic. In addition, the algorithm provides for statistical load balancing, not absolute load balancing. Therefore, since load balancing decisions are made on a conversation basis, it is possible that transmitted frames will not be equally distributed across all adapters in a team.

Implementers of Compaq Network Adapter Teaming can choose the appropriate algorithm for load balancing via the Compaq Teaming and Configuration Utility. This utility is a GUI used to configure Compaq Network Adapter Teaming and is installed when the Teaming driver is installed.

TLB and Layer 2 load balancing using MAC address

This algorithm makes load balancing decisions based on the destination MAC address of the frame being transmitted by the Teaming driver. The destination MAC address of the frame is the MAC address that belongs to the network device that will ultimately receive the frame. The Teaming driver utilizes the last three bits of the destination MAC address and assigns the frame to an adapter for transmission.

Since MAC addresses are in hexadecimal format, it is necessary to convert them to binary format. For example (see Table 1), a MAC address of 01-02-03-04-05-06 (hexadecimal) would be 0000 0001 – 0000 0010 – 0000 0011 – 0000 0100 – 0000 0101 – 0000 0110 in binary format. The Teaming driver load

balances based upon the last three bits (110) of the least significant byte (0000 0110 = 06) of the MAC address. Utilizing these three bits, the Teaming driver will consecutively assign destination MAC addresses to each functional network adapter in its team starting with 000 being assigned to network adapter 1, 001 being assigned to network adapter 2, and so on. Of course, how the MAC addresses are assigned depends on how many network adapters are in the TLB team and how many of those adapters are in a functional state.

Table 1. TLB Using MAC Addressing

Two Port Team:	**Three Port Team:**
MAC* Transmitting Adapter	MAC* Transmitting Adapter
000 = network adapter 1	000 = network adapter 1
001 = network adapter 2	001 = network adapter 2
010 = network adapter 1	010 = network adapter 3
011 = network adapter 2	011 = network adapter 1
100 = network adapter 1	100 = network adapter 2
101 = network adapter 2	101 = network adapter 3
110 = network adapter 1	110 = network adapter 1
111 = network adapter 2	111 = network adapter 2
Four Port Team:	**Five Port Team:**
MAC* Transmitting Adapter	MAC* Transmitting Adapter
000 = network adapter 1	000 = network adapter 1
001 = network adapter 2	001 = network adapter 2
010 = network adapter 3	010 = network adapter 3
011 = network adapter 4	011 = network adapter 4
100 = network adapter 1	100 = network adapter 5
101 = network adapter 2	101 = network adapter 1
110 = network adapter 3	110 = network adapter 2
111 = network adapter 4	111 = network adapter 3
	* MAC is the last three bits of the least significant byte.

Scenario 2

Taking the concepts reviewed in the section titled, "Scenarios of Network Addressing and Communication", this section describes how TLB MAC addressed- based load balancing functions.

Beginning at the point in Scenario 1 where Blue/1.1.1.2 transmits the PING Reply to Red/1.1.1.1, Blue must decide whether to use network adapter B or E. Blue's Teaming driver calculates using the MAC address of Red (A) because Red is the frame's destination. Because a hexadecimal "A" is equal to "1010" in binary, and the last three bits (010) are used to determine the transmitting network adapter (see (Table 2. 1 for a team with two network adapters), "010" is assigned to network adapter 1 (or the Primary Adapter). Therefore, when communicating with Red, Blue will always use the Primary Adapter to transmit frames.

If Blue transmits a frame to Yellow, the same calculation must be made. Yellow's MAC address is hexadecimal "F," which is equal to "1111" in binary. Blue's Teaming driver will again use the last three bits to determine which network adapter will transmit the frame. Referring to (Table 2. 1 for a team with

two network adapters, “111” is assigned to network adapter 2 (or the Secondary Adapter). Therefore, when communicating with Yellow, Blue will always use the Secondary Adapter to transmit frames.

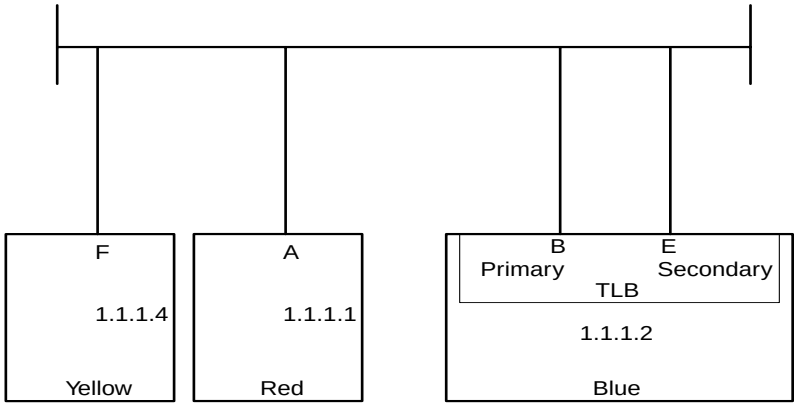

Figure 7. TLB Team Using MAC Address for Load Balancing Algorithm

TLB and Layer 3 load balancing using IP address

This algorithm makes load balancing decisions based on the destination IP address of the frame being transmitted by the Teaming driver. The frame’s destination IP address is that which belongs to the network device that will ultimately receive the frame. The Teaming driver utilizes the last three bits of the destination IP address to assign the frame to an adapter for transmission.

Because IP addresses are in decimal format, it is necessary to convert them to binary format. For example, an IP address of 1.2.3.4 (decimal) would be 0000 0001 . 0000 0010 . 0000 0011 . 0000 0100 in binary format. The Teaming driver only uses the last three bits (100) of the least significant byte (0000 0100 = 4) of the IP address. Utilizing these three bits, the Teaming driver will consecutively assign destination IP addresses to each functional network adapter in its team starting with 000 being assigned to network adapter 1, 001 being assigned to network adapter 2, and so on. Of course, how the IP addresses are assigned depends on the number of network adapters in the TLB team and how many of those adapters are in a functional state (see Table 2).

Table 2. TLB Using IP Addressing

Two Port Team:	**Three Port Team:**
IP* Transmitting Adapter 000 = network adapter 1 001 = network adapter 2 010 = network adapter 1 011 = network adapter 2 100 = network adapter 1 101 = network adapter 2 110 = network adapter 1 111 = network adapter 2	IP* Transmitting Adapter 000 = network adapter 1 001 = network adapter 2 010 = network adapter 3 011 = network adapter 1 100 = network adapter 2 101 = network adapter 3 110 = network adapter 1 111 = network adapter 2
Four Port Team:	**Five Port Team:**
IP* Transmitting Adapter 000 = network adapter 1 001 = network adapter 2 010 = network adapter 3 011 = network adapter 4 100 = network adapter 1 101 = network adapter 2 110 = network adapter 3 111 = network adapter 4	IP* Transmitting Adapter 000 = network adapter 1 001 = network adapter 2 010 = network adapter 3 011 = network adapter 4 100 = network adapter 5 101 = network adapter 1 110 = network adapter 2 111 = network adapter 3 * IP is the last three bits of the least significant byte.

Taking the concepts previously reviewed in Scenario 2 of the section titled, "Scenarios of Network Addressing and Communication", and Figure 8, this section describes how TLB IP addressed- based load balancing functions.

Beginning at the point in Scenario 2 where Blue/2.2.2.1 transmits the PING Reply to Red/1.1.1.1, Blue must decide whether to use network adapter B or E. Blue's Teaming driver calculates using the IP address of Red (1.1.1.1) because Red is the frame's destination. Because a decimal "1.1.1.1" is equal to "0000 0001 . 0000 0001 . 0000 0001 . 0000 0001" in binary, and the last three bits (001) are used to determine the transmitting network adapter (see Table 2. 2 for a 2 Network Adapter Team), "001" is assigned to network adapter 2 (or the Secondary Adapter). Therefore, when communicating with Red, Blue will always use the Secondary Adapter to transmit frames.

If Blue transmits a frame to Yellow, the same calculation must be made. Yellow's IP address is decimal "1.1.1.4" and equal to "0000 0001 . 0000 0001 . 0000 0001 . 0000 0100" in binary. Blue's Teaming driver will again use the last three bits to determine which network adapter will transmit the frame. Referring to Table 2 for a 2 Network adapter Team, "100" is assigned to network adapter 1 (or the Primary Adapter). Therefore, when communicating with Yellow, Blue will always use the Primary Adapter to transmit frames.

It is important to note that if an implementer uses the MAC address load balancing algorithm for the network in Figure 8, load balancing will not function as expected, and traffic will not be load balanced using all Teamed network adapters. Because Blue transmits all frames destined for Red and Yellow via Green (Blue's Gateway), Blue uses Green's Layer 2 address (MAC) as the frame's DESTINATION MAC ADDRESS but uses Red's and Yellow's Layer 3 addresses (IP) as the frame's DESTINATION IP ADDRESS. Blue never transmits frames directly to Red's or Yellow's MAC address because Blue is on a different Layer 2 network. Because Blue always transmits to Red and Yellow using Green's MAC address, the Teaming driver will assign all conversations with clients on Network 1.0.0.0 to the same network

adapter. When a Compaq Network Adapter Team needs to load balance traffic that traverses a Layer 3 device (Router), IP address based load balancing should be used.

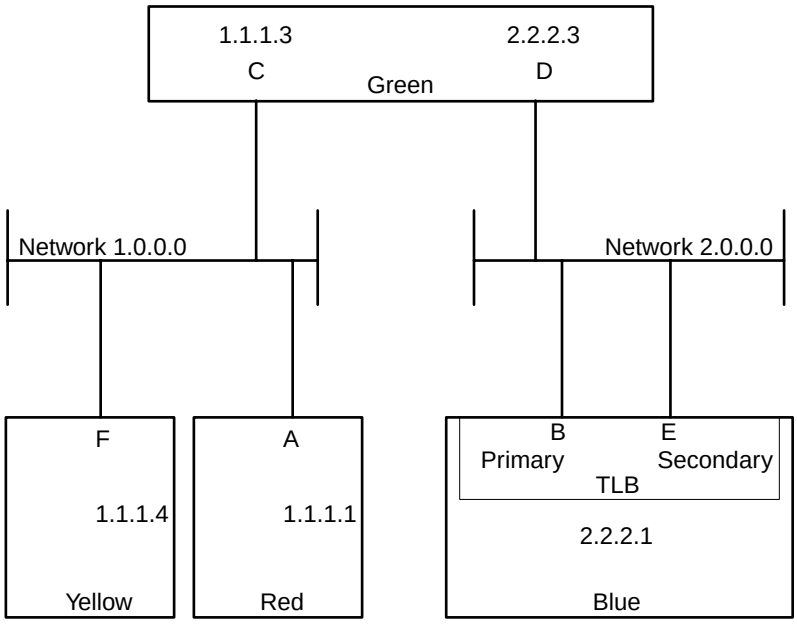

Figure 8. TLB Team Using IP Address for Load Balancing Algorithm

TLB Applications

TLB is deployed in environments that require fault tolerance and additional transmit throughput greater than the capacity of the Primary Adapter. TLB environments do not require receive throughput greater than the capacity of the Primary Adapter. (E.g., a database server whose primary role is that of transmitting data to clients. Receive throughput requirements may be much smaller than the transmit requirements.)

Recommended Configurations for a TLB Environment

The recommended configuration for a TLB environment is to have all members of the same TLB team attached to the same switch or hub. If switch redundancy is required (team members are attached to two different switches), then Compaq recommends that the switches be deployed with redundant links between them and Spanning Tree enabled on the ports that connect the switches.

Additionally, Compaq recommends that:

- Heartbeats are enabled (default) and the TLB team MAC address is not manually set to a locally administered address (LAA). A user should not LAA the MAC address of network adapters that are members of a team; otherwise Teaming may not function correctly.
- For TLB teams that communicate with network devices via a router, that the IP address-based load balancing algorithm be used.

- Spanning Tree be disabled on all switch ports to which a Compaq Network Adapter Team is attached. Cisco switches have a feature called Port Fast that is used to disable Spanning Tree on a port-by-port basis. If the Compaq-recommended configuration is followed (team members are attached to the same switch), disable Spanning Tree on all switch ports to which the teams are attached.

Switch-assisted Load Balancing (SLB)

Switch-assisted Load Balancing mode, formerly known as Fast EtherChannel mode (FEC) or Gigabit EtherChannel mode (GEC), incorporates all the features of NFT and TLB and adds the feature of load balancing of received traffic. In this mode, two to eight adapters may be teamed together as a single virtual network adapter. The load-balancing algorithm used in SLB allows for the load balancing of both the server's transmit and receive traffic (see Figure 9). Unlike TLB, which only load balances IP traffic, SLB load balances all traffic regardless of the Protocol.

Switch-assisted Load Balancing (SLB) is a Compaq term that refers to an industry standard technology for grouping multiple network adapters into one virtual network adapter and multiple switch ports one virtual switch port. Compaq has chosen to use the term "Switch-assisted Load Balancing" instead of adopting another vendor's terminology because Compaq's SLB technology works with multiple vendors' technology. Other compatible technologies from other vendors include: Cisco Fast EtherChannel (FEC)/Gigabit EtherChannel (GEC) (Static Mode Only), IEEE 802.3ad Link Aggregation (Static Mode only), Bay MultiLink Trunking, and Extreme Network's Load Sharing.

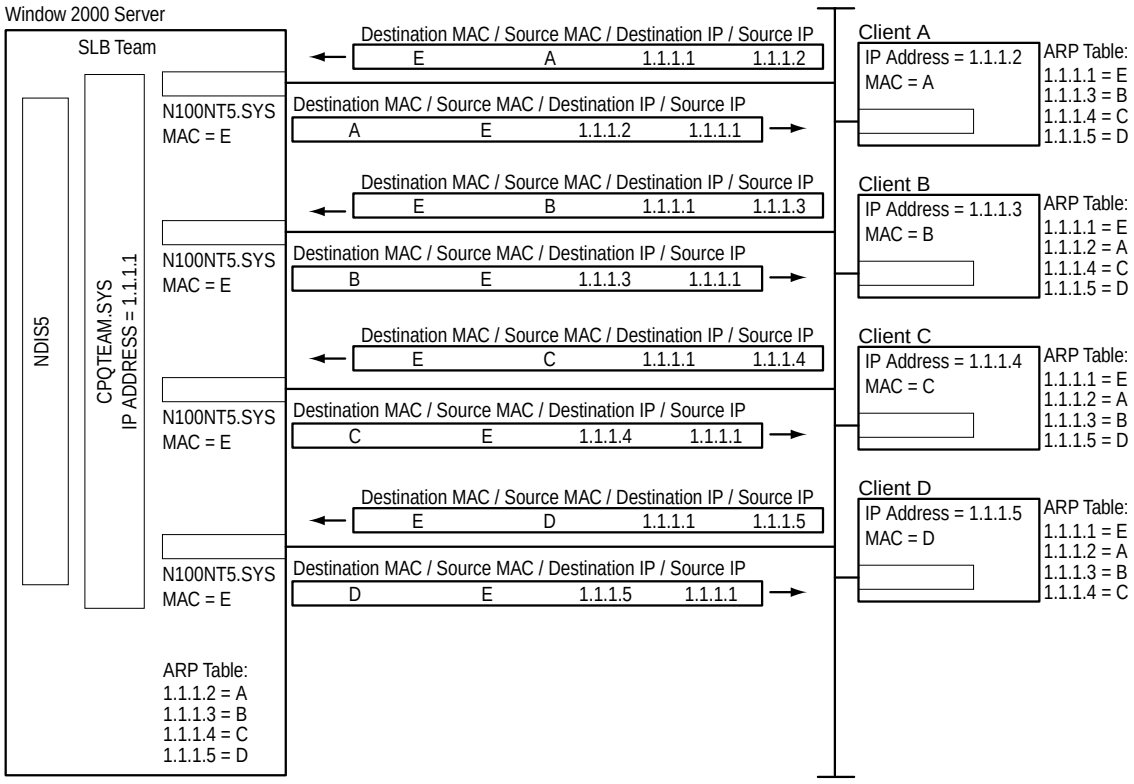

Figure 9. Overview of SLB Communication

Unlike NFT and TLB, SLB does not incorporate the concepts of Primary and Secondary Adapters within a team. All adapters within a team are considered equal and perform identical functions as long as the particular adapter is in a functioning state. The algorithm for Transmit load balancing used by SLB is identical to the algorithm used by TLB. Unlike TLB, SLB load balances all traffic regardless of the Protocol being used.

SLB and Layer 3 Load Balancing using IP Address

The algorithm for Transmit load balancing used by SLB is identical to the algorithm used by TLB. (See Section titled, "TLB and Layer 3 load balancing using IP address".)

Switch-assisted Load Balancing Receive Balancing Algorithm

The switch determines which load balancing algorithm is used to load balance receive traffic for an SLB team. An SLB team does not control which adapter in the team receives the incoming traffic. Only the switch can choose which adapter to use to send the traffic to the server. Therefore, please consult the switch manufacturer to determine the algorithm the switch uses.

Switch-assisted Load Balancing and Cisco's EtherChannel Technology

Teamed Cisco's Fast EtherChannel (FEC) and Gigabit EtherChannel (GEC) technology is a MAC layer (Layer 2) load balancing technology using two to eight network adapters grouped together as one logical network adapter. Depending on the specific load balancing algorithm used, FEC/GEC may not efficiently load balance traffic to network adapters.

FEC/GEC was originally designed as a switch-to-switch technology allowing two switches to increase the bandwidth between each other by aggregating multiple ports together as a single logical port for both transmits and receives. This is in contrast to Transmit Load Balancing (TLB) that only balances transmits. An algorithm had to be used that could statistically divide the traffic over each port in the FEC/GEC group in an attempt to divide it evenly.

There have been at least three algorithms that have been developed: source-based, destination-based, and XOR (see Table 3 for examples). The Source-based algorithm utilizes the last one or two bits (depending on the number of ports in the FEC/GEC group) of the source address in the packet. If the bit is 0, the first port is used. If the bit is 1, the second port is used. The Destination-based algorithm utilizes the last one or two bits (depending on the number of ports in the FEC/GEC group) of the destination address in the packet. If the bit is 0, the first port is used. If the bit is 1, the second port is used. The XOR algorithm utilizes the last one or two bits (depending on the number of ports in the FEC/GEC group) of the destination AND source addresses in the packet. The algorithm XORs the bits. If the result is 0, then the first port is used. If the result is 1, then the second port is used.

FEC/GEC has developed into not only a switch-to-switch technology but also a switch-to-node technology. In most cases, the node is a multi-homed server with network adapter drivers that support FEC/GEC. Problems can arise with switches using the destination-based algorithm when switch-to-node FEC/GEC is used. Because the destination address of the FEC/GEC node is always the same, the switch always sends traffic to that server on the same port. Because of this, receive traffic is not evenly distributed across all ports in the FEC/GEC group.

Table 3. Example of Algorithms

Preliminary Information:		
2 Port FEC/GEC group MAC address	= 00-00-00-00-00-**01**	(*Hexadecimal*)
Last Byte (**01**) in Binary	= 0000 **0001**	(*Binary*)
Client1 MAC address	= 00-00-00-00-00-**02**	(*Hexadecimal*)
Last Byte (**02**) in Binary	= 0000 **0010**	(*Binary*)
Client2 MAC address	= 00-00-00-00-00-**03**	(*Hexadecimal*)
Last Byte (**03**) in Binary	= 0000 **0011**	(*Binary*)
Packet 1 is a frame transmitted from Client 2 to the port FEC/GEC group. Packet 2 is a frame transmitted from Client 1 to the 2 port FEC/GEC group.		
DESTINATION-BASED ALGORITHM		
1. Packet 1 - Destination MAC address:		00-00-00-00-00-01
Last binary bit = 1 so frame is transmitted on port 2		
2. Packet 2 - Destination MAC address:		00-00-00-00-00-01
Last binary bit = 1 so frame is transmitted on port 2		
SOURCE-BASED ALGORITHM		
1. Packet 1 - Source Mac address:		00-00-00-00-00-03
Last binary bit = 1 so frame is transmitted on port 2		
2. Packet 2 - Source MAC address:		00-00-00-00-00-02
Last binary bit = 0 so frame is transmitted on port 1		
XOR ALGORITHM		
1. Packet 1 - Source Mac address:		00-00-00-00-00-03
Last binary bit = 1		
Packet 1 - Destination MAC address:		00-00-00-00-00-01
Last binary bit = 1		
(XOR result of binary bits 1 & 1 = 0 so frame is transmitted on port 1)		
2. Packet 2 - Source MAC address:		00-00-00-00-00-02
Last binary bit = 0		
Packet 2 - Destination MAC address:		00-00-00-00-00-01
Last binary bit = 1		
(XOR result of binary bits 0 & 1 = 1 so frame is transmitted on port 2)		

The effects of the destination-based algorithm do not indicate a fault in the network adapter drivers nor on the switch. Destination-based load balancing is considered a functional FEC/GEC algorithm because packets between switches may not always use the same destination or source addresses. Only single node-to-switch FEC/GEC uses the same destination address.

The algorithm used for load balancing has no effect on fault tolerance and fault tolerance will function the same in any implementation.

Some switches have the option to change the load balancing algorithm. In such cases, Compaq advises using the algorithms in this order of preference: XOR, source-based, destination-based.

Network Addressing and Communication using SLB

SLB functions identically to TLB (see section titled, "Network Addressing and Communication using TLB") except in its use of MAC addresses. Because SLB requires a switch capable of grouping multiple switch ports as a single switch port, and SLB Teaming uses the same MAC address on all network adapters in the same team, this does not violate IEEE standards since the switch is fully aware of the port groupings and expects that all network adapters will transmit using the same MAC address.

SLB Applications

Switch-assisted Load Balancing is deployed in environments that require fault tolerance and additional transmit and receive throughput greater than the capacity of the Primary Adapter and that have a switch capable of providing, and configured to provide load balancing assistance (e.g., a backup server that requires additional receive throughput for backing up other servers and clients).

Recommended Configurations for an SLB Environment

Because SLB is highly dependent on the switch vendor's implementation of port bonding/aggregation/teaming, Compaq recommends that implementers thoroughly understand the configuration guidelines set by the switch vendor. Compaq's SLB technology has been designed to allow for flexibility. Therefore, the Compaq Teaming And Configuration GUI may allow configuration of an SLB team that will not work correctly with a particular vendor's switch.

Compaq recommends that the IP address-based load balancing algorithm be used for SLB teams that will communicate with network devices via a router.

Compaq also recommends disabling Spanning Tree Algorithm on all switch ports to which a Compaq Network Adapter Team is attached. Cisco switches have a feature called Port Fast that is used to disable Spanning Tree on a port-by-port basis.

Network Adapter Fail Over

NFT and Network Adapter Failure Recovery

There are three Fail Over modes available for NFT teams: Manual, Fail On Fault, and Preferred Primary.

Manual Mode

This mode for NFT does not provide any type of fault tolerance. Manual mode is used to disable the fail over feature of NFT, and for user-initiated failovers (manual failovers). When set, Manual mode prevents an NFT team from automatically failing over during events that normally cause a fail over (e.g., if a cable were pulled on a Primary Adapter of an NFT team). Manual mode is normally used for troubleshooting

purposes (e.g., using an analyzer to take an inline network trace). Compaq does not recommend using this mode as the permanent mode in a normal operating environment.

Fail On Fault Mode

The second mode available for NFT is Fail On Fault. In this mode, an NFT team will fail over from the Primary Adapter to an operational Secondary Adapter whenever a failover event occurs (see Section 0 Fail Over Events) on the Primary Adapter. When the failover occurs, the two adapters swap MAC addresses so the Team remains known to the network as the same MAC address. The new Primary Adapter is considered just as functional as the old Primary Adapter. If the old Primary Adapter is restored, it becomes a Secondary Adapter for the team but no MAC address changes are made unless there is another failover event on the Primary Adapter.

Preferred Primary Mode

The last mode available for NFT is Preferred Primary mode. When choosing Preferred Primary mode, the operator is presented with a drop down box to select the "Preferred Primary Adapter". The operator should choose the adapter that, for a particular reason, is best suited to be the Primary Adapter. For instance, if an NFT team were to be created using a Gigabit adapter and a 10/100 adapter, an operator would choose the Gigabit adapter as the Preferred Primary Adapter, because of the increased bandwidth available with a Gigabit adapter.

When an adapter is chosen as the Preferred Primary Adapter, it will be used as the Primary Adapter whenever it is in an operational state. If the Preferred Primary Adapter experiences a failover event, the NFT team fails over to a Secondary Adapter. If the Preferred Primary Adapter is restored, the team will then fail back to the Preferred Primary Adapter. Essentially, the team will fail over twice even though only one error occurred. The second fail over is to make the restored Preferred Primary Adapter the team's Primary Adapter once again. A fail over back to the Preferred Primary is more specifically referred to as a fail back.

Note: Failures of Secondary Adapters do not trigger any type of recovery since Secondary Adapters are already in standby mode. The only consequence of a failed Secondary Adapter is the possibility of the Primary Adapter failing and the Team becoming unavailable to the network because both adapters in the Team are in a failed state. If there are three or more network adapters in a team and two adapters fail, the Team is still available via the third adapter.

TLB and Network Adapter Failure Recovery

With TLB, the recovery mechanism provided is very similar to the NFT fail over mode discussed in section titled, "Fail On Fault". In a 2-port TLB team, the primary adapter receives all data frames, while the Secondary Adapter receives only heartbeat frames. Both adapters are capable of transmitting data frames. In the event of a failover, the Secondary Adapter becomes the Primary Adapter and assumes the MAC address of the team. In effect, the two adapters swap MAC addresses. The new Primary Adapter now receives and transmits all data frames. If the old Primary Adapter is restored, it becomes a Secondary Adapter for the team. It will now only receive heartbeat frames and be capable of transmitting data frames. If a Secondary Adapter fails in a two-port team, the data frames being load balanced by the adapter are transmitted by the Primary Adapter. If the Secondary Adapter is restored, it remains secondary, and the team will resume load balancing data frames on that adapter. No MAC address changes are made when a Secondary Adapter fails or is restored.

SLB and Network Adapter Failure Recovery

With SLB, the recovery mechanism is somewhat different than those discussed previously. All members of the team transmit and receive frames with the same MAC Address and there are is no concept of a primary or Secondary Adapter as there is in NFT and TLB. With SLB, there are no heartbeat frames, and consequently no heartbeat failovers. In a two-port SLB team, all members are capable of receiving data frames (based on the switch's load balancing algorithm), and transmitting data frames (based on the Teaming Driver's load balancing algorithm). In the event of a fail over, all transmit traffic is redistributed among the working adapters. After a failover event in a two-port team, only one adapter is currently working, so all transmit traffic is sent using it. All receive traffic is determined by the switch, which should detect that only one adapter is working. If the failed adapter is restored, all transmit traffic is once again load balanced among all adapters.

All receive traffic is still determined by the switch algorithm, which should detect that both adapters are functional. If the switch sends a frame destined for the team MAC address to any of the "operational" adapters in the team, the adapter will receive it. The Compaq Network Adapter Teaming driver does not control frames received, but only load balances the transmit traffic. All protocols are load balanced, not just IP. Remember that the Teaming driver load balances network conversations per teamed adapter, and therefore it is possible that transmit traffic being sent out a particular adapter to a particular device could change ports after a failure on another adapter in the team. This could occur because the Teaming driver algorithm redistributes transmit traffic among working adapters every time the state of any member of the team changes.

Unlike NFT and TLB, a failure on any adapter within the same SLB team has the same ramifications as a failure on any other adapter since all adapters are considered equal.

Fail Over Events

Link Loss

When a network adapter is a member of a Team and loses physical link (i.e., link light is lost), the Teaming driver disables that adapter in the team. If this adapter is in use by the Team, the Team recovers from the failure based on the adapter's role in the Team (Primary or Secondary) and the Team's mode (NFT, TLB, or SLB). (See sections 0, 0, and 0.)

Heartbeat Failures

The use of Heartbeat frames for network adapter failovers was designed to detect network adapter communication failure even in cases when it maintained physical link. Special Heartbeat frames (see Section 0 Heartbeats) are transmitted and received by teamed network adapters to validate the transmit and receive paths of each adapter. When a Heartbeat frame is transmitted by one teamed adapter but not received by another teamed adapter, the Teaming driver assumes that one of the adapters is having a communications problem. If the Primary Adapter in an NFT or TLB team experiences a failover event, a failover may occur.

Heartbeat frames were not designed to monitor for other networking failures such as loss of connectivity between switches, loss of client connectivity.

Heartbeats

Heartbeats are special frames that Compaq's Network Adapter Teaming uses for validating team member network connections and for notifying other network equipment of MAC address changes as a result of fail

over events. See the section titled, " Heartbeat Functionality and Timers" for a complete description of the different heartbeat types.

Heartbeat frames contain only Layer 2 addresses (refer to the section titled, " Heartbeat Frame Format" for the complete frame format of a heartbeat) and do not contain any Layer 3 addresses. This means that heartbeat frames are not routable. In other words, heartbeat frames will not be routed (by a router) between team members if the team members are on two different Layer 2 networks joined by a Layer 3 device (router). If the heartbeat frames aren't delivered between team members, then erroneous failovers may occur.

Heartbeat Frame Format

	802.2 Frame Format
Destination MAC address:	03-00-C7-00-00-EE (Multicast)
Source MAC address:	<current MAC address of transmitting adapter>
Frame Length:	66 bytes
DSAP:	AA
SSAP:	AA
SNAP Type:	Unnumbered, TEST
Data:	63 bytes of insignificant data

Heartbeat Functionality and Timers

There are three main functions of the heartbeat frame used for Compaq Network Adapter Teaming: Receive Validation, Transmit Validation, and Switch MAC Table Updates. The same heartbeat frame format is used for all heartbeat functions and they are indistinguishable from the network's perspective.

Compaq Network Adapter Teaming utilizes a mechanism on a timer for checking network connectivity with heartbeats. This timer is called the Heartbeat Timer Interval and is available for custom tuning on the Settings Tab of the Properties Menu of a team in the Compaq Teaming and Configuration GUI. The default value of this timer is 3000 (milliseconds), and the valid range is 3000 milliseconds to 60000 milliseconds (3 seconds to 60 seconds).

The terms HSM and TVSM used in the following three sections are used in this document only for the purpose of explaining heartbeat functionality.

Receive Validate Heartbeats

Teaming uses the Heartbeat Timer Interval to know how often to validate receives. The driver executes a Heartbeat State Machine (HSM) every three seconds (default). If no receives have occurred in that three second interval, the Adapter Heartbeat status increases by 1. Once the Adapter Heartbeat status increases to 3 for a Primary Adapter, a heartbeat is transmitted by all Secondary Adapters. If one of the Secondary Adapters has degraded to 3, then the Primary Adapter transmits a heartbeat. 3 seconds multiplied by 3 state transitions = 9 seconds. Therefore, a Heartbeat frame will be transmitted from the Primary Adapter to a Secondary Adapter or from all Secondary Adapters (note that this is plural in a team of three or more members) to the Primary Adapter every nine seconds if there is no receive activity on a particular team member.

Switch MAC Table Update with Team Address Heartbeat

The Primary Adapter must ensure that the switch has the team's MAC address on the correct port, and that the switch does not timeout the Team's MAC if the Primary Adapter has not transmitted for a while (only an issue with TLB since TLB transmits on the Secondary Adapters and it may happen that the Primary Adapter only gets used for receives). So, the Primary Adapter transmits a heartbeat on every pass of the HSM (every three seconds by default).

Transmit Validation Heartbeat

All teamed adapters (Primary and Secondary) pass through a Transmit Validation State Machine (TVSM) once every three times through the HSM. This TVSM checks to see if each adapter has successfully transmitted since the last pass through the TVSM. If not, then the adapter's internal status is incremented by 1 (starts at 0). Once the adapter's internal status reaches 3 (4 transitions) without transmitting something, that adapter will transmit a heartbeat. Therefore, a Transmit Validation Heartbeat may be transmitted once every 36 seconds (if heartbeat interval timer is set to default of 3000ms) on an adapter that has not transmitted anything. HSM (executes every 3 seconds) x TVSM (executes every 3 times through HSM) x adapter status transitioning 4 times: (3 x 3 x 4 = 36 seconds).

Compaq Network Adapter Teaming and Advanced Networking Features

Checksum Offloading

There are three features, Receive TCP Checksum Offloading, Transmit TCP Checksum Offloading, and Transmit IP Checksum Offloading that are supported on Compaq NC Series Gigabit adapters. Compaq Network Adapter Teaming does not support the Offloading advanced features. These features are disabled when the network adapters are installed into a team. Future enhancements to the software will provide support for these features. Offloading features on NC Series Gigabit Adapters are supported only in non-teamed environments.

802.1p QoS Tagging

Compaq Network Adapter Teaming supports the 802.1p QoS advanced feature. If any one team member does not support 802.1p QOS, then it will be disabled on all team members. If all team members support 802.1p QoS, and any one team member has it enabled, then it will be enabled on all team members. 801.2p QoS is supported on NC Series Fast Ethernet/Gigabit Adapters.

Gigabit Jumbo Frames

Compaq Network Adapter Teaming supports the Maximum Frame Size advanced feature (Jumbo Frames). If one or more team members do not support Jumbo Frames, then it will be disabled on all team members. Disabled is equal to 1514 Bytes. If all team members support Jumbo Frames, the teaming driver will use the lowest common denominator frame size and set that value for all team members. For example, if there are two teamed adapters, one configured for 4088 Bytes and the other for 9014 Bytes, the Maximum Frame Size feature is set to 4088 Bytes. Note that Maximum Frame Size greater than 1514 would constitute Jumbo Frames being enabled. Jumbo Frames are supported on NC Series Gigabit Adapters only. In addition, the specified Maximum Frame Size in Compaq Network Adapter Teaming does not include the four-byte Cyclic Redundancy Check (CRC) portion of an Ethernet frame. Some switch settings do include the CRC portion in their Jumbo Frame Size configuration. Therefore, it may be necessary to increase the switch's Jumbo Frame Size setting by four bytes in order for Jumbo Frames to work correctly.

Network Scenario Considerations

NFT/TLB Team Split Across Switches

Compaq Network Adapter Teaming is designed for network adapter fault tolerance. However, some System or Network Administrators may deploy Compaq Network Adapter Teaming with switch fault tolerance in mind (see Figure 10). A thorough understanding of the heartbeat process and network adapter failover mechanisms is necessary before implementing NFT or TLB in this type of environment.

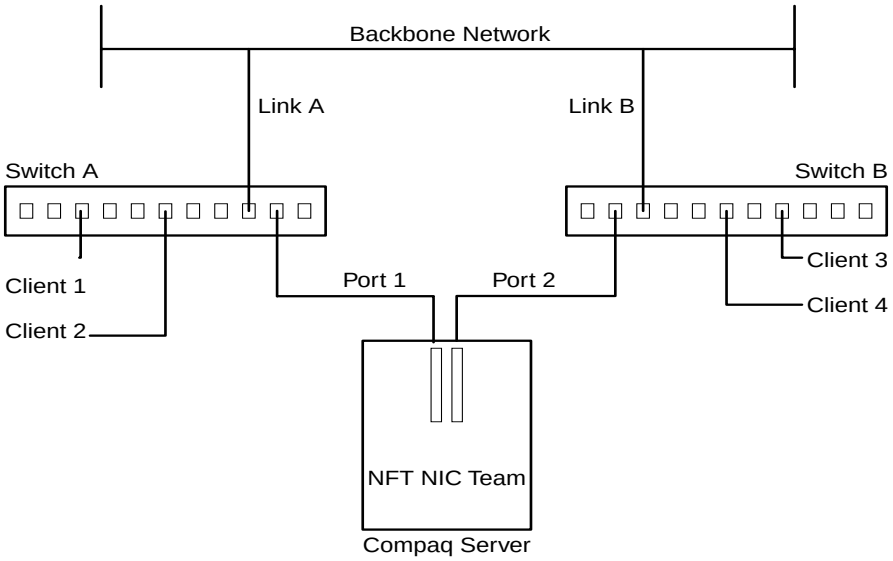

Figure 10. NFT/TLB Team Split Across Switches

For example, in Figure 10, the Compaq Network Adapter Team is attached to two switches, A and B. Both switches are uplinked to a core/backbone network. Network adapter 1 (NIC 1) is the Primary Adapter in this NFT team. Network adapter 2 (NIC 2) is the Secondary Adapter. Because this is an NFT team, the Secondary Adapter is in standby mode. A typical team member failure is caused by link loss on that team member. Immediately, the Teaming driver distinguishes the problem. If the link loss is on the Primary Adapter (NIC 1), then the Teaming driver will fail over to NIC 2 if NIC 2 is functioning (heartbeats have been successful).

Another type of failure can also occur in this scenario. Instead of one of the Team adapters losing link, the link from Switch A or B to the backbone network could be lost (Link A or B). This type of failure would isolate either switch from the backbone network and from the other switch. It would also cause the Compaq Network Adapter Team to enter Heartbeat failure mode because Heartbeat packets would not be successful between NIC 1 and NIC 2. A problem can arise when the Teaming driver makes a failover decision at this point because the Teaming driver has no way to determine where the failure occurred in the network. If the failure occurred at Link A, the best decision the Teaming driver could make is to fail over to NIC 2 and let NIC 2 be the new Primary Adapter. This decision would give server access to the most clients, Clients 5, 4 and 3; and would isolate the fewest number of clients, Clients 1 and 2. However, because the Teaming driver cannot determine which link failed, the Teaming driver will fail over to the Secondary Adapter every time a heartbeat failure occurs as long as the Secondary Adapter is in an operational state. This means that if Link B were to fail, the Teaming driver would fail over to NIC 2 (the Secondary Adapter). This is not the optimal decision since it isolates more clients. Clients 3 and 4 have access but Clients 1, 2, and 5 are isolated.

One can improve this issue by implementing redundant links between the switches and deploying Spanning Tree Protocol (STP) on the switched LAN (see Figure 11). In this diagram, Switch A and Switch B have been redundantly connected, both directly and via the backbone network. In addition, STP has been implemented and has put the redundant link in standby mode (STP Block) until a failure occurs on the network that segments Switch A from Switch B. When the failure occurs, the link between the switches is made operational and the Compaq Network Adapter Team is able to communicate with all clients without isolating any.

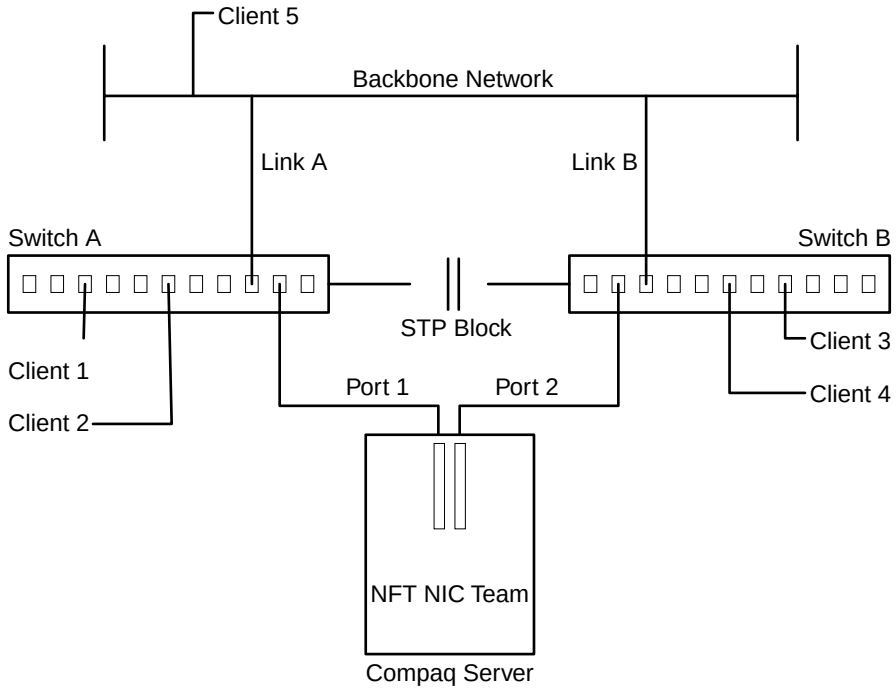

Figure 11. Deploying Spanning Tree Protocol on a Switched LAN

NFT/Preferred Primary Team Split Across Switches

Using the same scenario as in Figure 10, an NFT Team in Preferred Primary mode will behave differently than a TLB Team or NFT Team in Fail-On-Fault mode. In Preferred Primary mode, an NFT Team will use the adapter designated as Preferred Primary as long as the adapter is functioning. If Link A or Link B fails in the above network scenario, loss of Heartbeat frames will initiate an adapter failover and NIC 2 will take over as the new Primary Adapter. However, NIC 1 is elevated to an operational state if it receives any frame (unicast, multicast or broadcast). Therefore, even though a heartbeat failure just caused a failover to NIC 2, the Teaming driver will fail back to NIC 1 as soon as it receives a frame. Once NIC 1 becomes the Primary Adapter again, the Team may immediately enter a heartbeat failure state and fail over to NIC 2. Once again, as soon as NIC 1 receives any type of frame, the Team will fail back to NIC 1. This can cause an endless loop of failovers.

Layer 3 Routing Of Load Balanced Traffic

Special consideration should be given when choosing between the MAC Address-based and IP Address-based load balancing algorithms in an environment where the server and clients are separated by a Layer 3 device, such as a router. In such an environment, the server must communicate with the clients via the router (set as its default gateway). When communicating with the clients, the server sends all traffic to the

router, which then sends the traffic to the clients. If MAC Address- based load balancing is selected, all traffic destined for clients is transmitted using the same network adapter in the load balancing team and is not load balanced. This occurs because the server must use the router's MAC address as the Layer 2 address in every frame while it uses the client's IP address as the Layer 3 address in the same frame. Because the router's MAC address is used in every frame, the MAC address-based load balancing algorithm chooses the same adapter for all traffic. Instead, choose the IP address-based load balancing algorithm, and load balancing will be based on the address of the clients (which varies) and not on the router (which is the same). In hybrid environments in which the server is communicating with clients on its own network, as well as clients on the other side of a router, Compaq recommends using IP-based load balancing. See Sections 0 TLB and Layer 3 load balancing using IP address, and 0 SLB and Layer 3 Load Balancing using IP Address for a more detailed explanation.

TLB and Layer 3 Switching

There are two main types of switches used for networking, Layer 2 switches and Layer 3 switches. Layer 2 switches make switching decisions based on the destination MAC address of the frame being transmitted. Layer 3 switches make decisions based on the Layer 3 address of the frame being transmitted.

When using TLB on a Layer 3 switch, implementers may encounter a problem called "switch thrashing" or high switch CPU utilization. The problem occurs because the switch sees the same source IP address on multiple ports (number of ports depends on the size of the team). The optimal solution for this issue is to implement a different mode of teaming, SLB, if the switch supports it. With SLB, the switch is aware of the network adapter team and does not have a problem with the use of a single IP address on more than one switch port.

Load Balancing and IPX Traffic

When using Compaq Network Adapter Teaming in Transmit Load Balancing (TLB) mode, frames destined for clients will be load balanced over each network adapter when transmitting. A unique Layer 2 address is used for each network adapter but the same Layer 3 address is used by all network adapters. Most protocols work perfectly in this environment; however, some IPX clients require that the Layer 2 address remain constant for any host using a particular Layer 3 address. Because of this, TLB mode does not load balance IPX traffic. IPX traffic is always transmitted out of the Primary Adapter and, in order to avoid any potential problems with other IPX clients, always uses only one Layer 2 address.

If load balancing of IPX traffic is required, a different mode of load balancing called SLB (Switch-assisted Load Balancing) should be used. With SLB, all traffic is transmitted from each network adapter using the same Layer 2 and Layer 3 addresses. Because SLB uses the same Layer 2 address, IPX traffic will be load balanced without affecting the clients.

Load Balancing and AppleTalk Traffic

When using Compaq Network Adapter Teaming in Transmit Load Balancing (TLB) mode, frames destined for clients will be load balanced over each network adapter when transmitting. A unique Layer 2 address is used for each network adapter but the same Layer 3 address is used by all network adapters. Most protocols work perfectly in this environment; however, some AppleTalk clients require that the Layer 2 address remain constant for any host using a particular Layer 3 Address. Because of this, TLB mode does not load balance AppleTalk traffic. AppleTalk traffic is always transmitted out of the Primary Adapter and, in order to avoid any potential problems with other AppleTalk clients, always uses only one OSI Layer 2 address.

If load balancing of AppleTalk traffic is required, a different mode of load balancing called SLB (Switch-assisted Load Balancing) should be used. With SLB, all traffic is transmitted from each network adapter using the same Layer 2 and Layer 3 addresses. Because SLB uses the same Layer 2 address, AppleTalk traffic will be load balanced without affecting the clients.

Load Balancing and SNA Traffic

When using Compaq Network Adapter Teaming in Transmit Load Balancing (TLB) mode, frames destined for clients will be load balanced over each network adapter when transmitting. A unique Layer 2 address is used for each network adapter but the same Layer 3 address is used by all network adapters. Most protocols work perfectly in this environment; however, SNA (native) requires that the Layer 2 address remain constant. Because of this, TLB mode does not load balance SNA traffic. SNA traffic is always transmitted out of the primary network adapter and, in order to avoid any potential problems with other SNA clients, always uses only one OSI Layer 2 address.

If load balancing of SNA traffic is necessary, a different mode of load balancing called SLB (Switch-assisted Load Balancing) can be implemented. With SLB, all traffic is transmitted from each network adapter using the same Layer 2 and Layer 3 addresses. Because SLB uses the same Layer 2 address, SNA will be load balanced without affecting the clients.

Teaming Feature Matrix

Teaming Type	NFT	TLB	SLB
Number of adapters supported per team	2-8	2-8	2-8
Supports Fault Tolerance	X	X	X
Supports Transmit Load Balancing		X	X
Supports Receive Load Balancing			X
Requires a switch that supports a compatible form of load balancing			X
Can connect a single team to more than one switch for switch redundancy (must be same broadcast domain)	X	X	Switch dependent
Utilizes heartbeats for network integrity checks	X	X	
Can team adapters that do not support a common speed	X		
Can team adapters operating at different speeds as long as the adapters support a common speed	X	X	X
Can team adapters of different media	X	X	X
Maximum theoretical transmit/receive throughput (in Mbps) with maximum number of 100 Mbps adapters	100/100	800/100	800/800
Maximum theoretical transmit/receive throughput (in Mbps) with maximum number of 1000 Mbps adapters	1000/1000	8000/1000	8000/8000
Load balances TCP/IP		X	X
Load balances IPX/SPX			X
Load balances SNA			X
Load balances AppleTalk			X
Supports load balancing by destination IP address		X	X
All adapters within a team utilize the same MAC address on the network			X
All adapters within a team utilize the same IP address on the network	X	X	X

Definitions and Acronyms

ALB	**A**daptive **L**oad **B**alancing. See Transmit Load Balancing (TLB).
ARP	**A**ddress **R**esolution **P**rotocol. A protocol used to determine a MAC address from an IP address.
BIA	**B**urned **I**n MAC **A**ddress. The Layer 2 address that is permanently assigned to a piece of hardware by the vendor.
Broadcast domain	Set of all devices that will receive Layer 2 broadcast frames originating from any device within the set. Broadcast domains are typically bounded by routers since routers do not typically forward Layer 2 broadcast frames.
Byte	Eight bits
Collision domain	A single Ethernet network segment in which there will be a collision if two computers attached to the system transmit simultaneously.
FEC	**F**ast **E**ther**C**hannel. A method of load balancing transmit and receive traffic across multiple Fast Ethernet connections (100 Mbps) between two devices. Developed by Cisco Systems. See SLB.
GEC	**G**igabit **E**ther**C**hannel. A method of load balancing both transmit and receive traffic across multiple Gigabit Ethernet (1000 Mbps) connections between two devices. Developed by Cisco Systems. See SLB.
GUI	**G**raphical **U**ser **I**nterface.
IEEE	**I**nstitute of **E**lectrical and **E**lectronics **E**ngineers. A standards body for, among other things, network communications and protocols.
LAA	**L**ocally **A**dministered **A**ddress. A temporary Layer 2 address that is manually assigned to a piece of hardware.
Layer 2	The second layer of the OSI model, the Data Link Layer. A Layer 2 address is the same as a MAC (Media Access Control) address.
Layer 3	The third layer of the OSI model, the Network Layer. A Layer 3 address refers to a protocol address such as an IP or IPX address.
MAC address	**M**edia **A**ccess **C**ontrol address. With Ethernet, this refers to the 6 byte (48 bit) address that is unique to every Ethernet device
Multi-homed	

	networks.
NDIS	**N**etwork **D**river **I**nterface **S**pecification. Simplified, it is the interface between a network adapter and Microsoft's protocol stack.
NFT	**N**etwork **F**ault **T**olerance. A team of network adapters that transmits and receives on only one adapter with all other adapters in standby.
OSI Model	**O**pen **S**ystems **I**nterconnect Model. The seven layer model developed by the International Standards Organization that outlines the mechanisms used by networked devices to communicate with each other.
PING	A type of packet used to validate connectivity with another network device. The packet asks another network device to respond to validate connectivity, a kind of "echo". PING packets for IP are accomplished using the ICMP protocol.
SLB	**S**witch-assisted **L**oad **B**alancing. Also known as FEC/GEC. A team of network adapters that load balances transmits and receives on all adapters.
STA	**S**panning **T**ree **A**lgorithm [(IEEE 802.1D)]
Switch MAC Table	A list of MAC addresses and associated ports that is used by a switch to transfer frames between attached devices.
TLB	**T**ransmit **L**oad **B**alancing. Also known as Adaptive Load Balancing (ALB). A team of network adapters that receives on one adapter but load balances transmitted IP traffic on all adapters. Other protocol traffic is transmitted by a single adapter.

Technical Support

To contact a Compaq engineer regarding issues with Compaq Network Adapter Teaming, please call 1-800-652-6672 (1-800-OK COMPAQ) or 1-800-386-2172 (Compaq Server and Networking Support). To speak with the appropriate support group, select the call routing options for Compaq Server Networking, or Compaq ProLiant Networking.

For online assistance, Compaq's Web based Support Forum is located at: http://www.compaq.com/support/.

For driver updates to Compaq Network Adapter Teaming, please visit Compaq's Network Adapter Driver site at: http://www.compaq.com/support/files/networking/nics/index.html.