

April 1999
ECG062/0299

Prepared by High Availability Group

Compaq Computer Corporation

Contents

Introduction	3
Communication points of failure	4
Supported Non-Fault Tolerant Architectures	5
Supported Fault Tolerant Architecture	6
Teaming Configuration for Higher Availability	8
Hardware and Software Requirements	8
Teaming Configuration of Ethernet Interconnects for Compaq Fast Ethernet NC Series NIC cards	9
Option 1. NIC Fault Tolerance (NFT) Configuration	9
Option 2. Adaptive Load Balancing (ALB) Configuration	9
Steps for Configuring NICs Utilizing Teaming	10
Verifying Client LAN and Cluster Interconnect IP Addresses	16

Supported Ethernet Interconnects for Compaq Parallel Database Clusters using Oracle Parallel Server

Abstract:

This Technical Guide reviews the supported Ethernet client LAN and cluster interconnect options for the Compaq Parallel Database Clusters for Oracle Parallel Server. Both redundant and non-redundant configurations are discussed. Hardware and software requirements along with limitations are detailed. The recommended cluster interconnect architecture is a fully redundant, and therefore highly available configuration.

Notice

The information in this publication is subject to change without notice and is provided “AS IS” WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, registered United States Patent and Trademark Office.

Netelligent, Armada, Cruiser, Concerto, QuickChoice, ProSignia, Systempro/XL, Net1, LTE Elite, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, Wonder Tools logo in black/white and color, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows, Windows NT, Windows NT Server and Workstation, Microsoft SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetWare and Novell are registered trademarks and intraNetWare, NDS, and Novell Directory Services are trademarks of Novell, Inc.

Pentium is a registered trademark of Intel Corporation.

Oracle8 is a trademark and Oracle is a registered trademark of Oracle Corporation.

Copyright ©1999 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

First Edition (April 1999)

Document Number ECG062/0299

Introduction

This Technical Guide is intended to be a supplement to the Order and Configuration Guides, Administrator Guides, and Certification Matrix for the Compaq Parallel Database Clusters for Oracle Parallel Server (OPS). This Guide reviews the supported Ethernet architectures for the client LAN and cluster interconnect (sometimes known as SAN) in an OPS environment. It does not provide detailed cluster network configuration information, nor does it address the ServerNet interconnect options available. Please refer to your Order and Configuration Guide and Administrator Guide for the configuration and implementation of these architectures.

The client LAN and cluster interconnect used in a Compaq Parallel Database Cluster environment can be a combination of redundant or non-redundant architectures. This document begins with a brief discussion of the non-redundant architecture for both the client LAN and cluster interconnect. Limitations and requirements are outlined.

Two options are detailed for implementation of fully redundant client LAN and cluster interconnect connections - NIC Fault Tolerant (NFT) and Adaptive Load Balancing (ALB). The ALB option is recommended for the best performance and availability of your cluster. Both configurations are fully tested and supported. These architectures will be defined for hardware, software and configuration requirements. These are the recommended configurations for the Compaq Parallel Database Clusters for OPS.

Communication points of failure

Several components make up the physical network of the client LAN and cluster interconnect communication paths. The failure of any one of these components renders the entire path inoperable. Unless redundancy has been designed into the communication paths, a component failure can cause a complete disruption of access to certain cluster resources.

The following four hardware items are the primary points of possible failure.

- A port on a dual-ported network controller (LAN or cluster interconnect)
- An entire network controller (LAN, cluster interconnect or a combination of both)
- A network cable
- A port on a switch
- A switch

An integration of hardware and software technologies can create redundancy to maintain a robust cluster, thus supplying defenses against points of failure in the architecture. This increases both the resiliency of cluster communications, and the overall availability of cluster applications and data. To protect against any communication disruption, Compaq highly recommends you implement the fully redundant communications paths specified in this document. The following are the redundancy characteristics of both the cluster interconnect and the client LAN connections in the fully redundant configurations recommended in the following sections.

The cluster interconnect remains available if:

a network controller port used for the primary interconnect path fails.

a network controller used for the primary interconnect path fails.

the network cable(s) used for the primary interconnect path fails.

the port(s) on the switch used for the primary interconnect path fails.

the switch used for the primary interconnect path fails.

Point(s) of Failure: None, if all client connections to the cluster interconnect switch are redundant.

The client LAN remains available if:

a network controller port used for the primary client LAN path fails.

a network controller used for the primary client LAN path fails.

the network cable(s) used for the primary client LAN path fails.

the port(s) on the switch used for the primary client LAN path fails.

the switch used for the primary interconnect path fails.

Point(s) of Failure: None, if all client connections to the client LAN switch are redundant.

Supported Non-Fault Tolerant Architectures

A non-fault tolerant architecture means that your cluster will have at least one point of failure in the communications network. Certain non-fault tolerant Ethernet architectures are supported.

For two node configurations, the use of a simple crossover cable is one of the supported options. See Figure 1 for an illustration of this option.

NOTE: Redundant crossover cables are NOT supported.

A single path connection through a hub or switch is also a supported two node non-fault tolerant configuration.

For more than two node configurations, a switch is required. (It doesn't necessarily need to be fully redundant to be supported.)

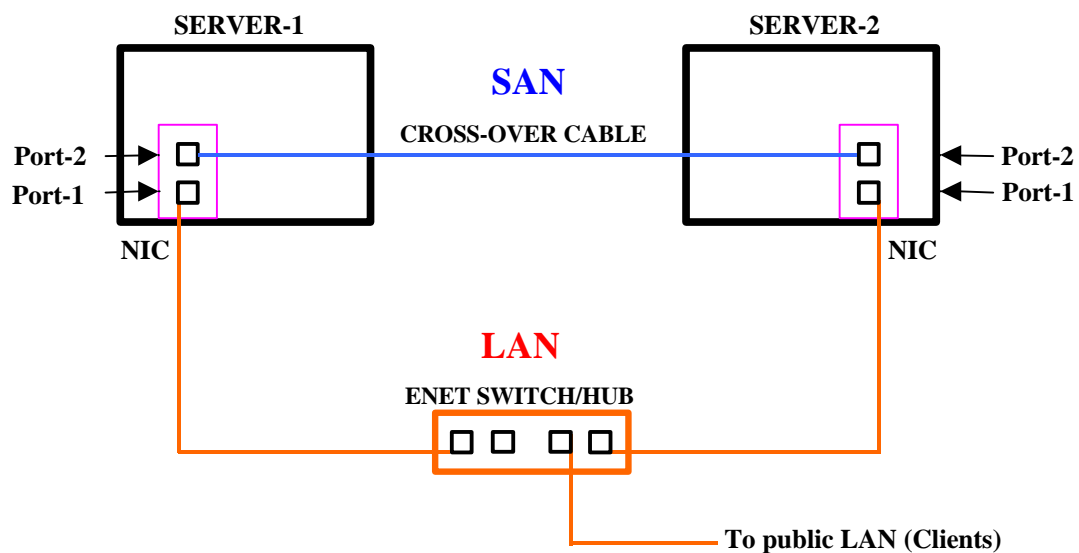


Figure 1. Example of a non-fault tolerant, 2-node configuration using a cross-over cable

Supported Fault Tolerant Architecture

A fault tolerant architecture consists of combining Microsoft NT and Compaq software with Compaq Fast Ethernet Network Interface Controllers (NICs) in the Compaq NC Series, to create redundancy. With this configuration, two NICs can be configured for a fault tolerant network with a path for primary and a path for secondary or backup communication on separate NICs (either dual or single port), for each node. This path/port teaming feature is enabled with a Compaq utility in the Microsoft Windows NT 4.0 network control panel.

Several configurations are supported for the cluster interconnect and the client LAN. In this document, only the most robust configuration that will provide both the highest availability and performance is recommended (See Figure 2). (This document will address redundancy in both the LAN and cluster interconnect. It is not required that both networks be fully redundant to be supported, any combination will work as well.)

The following configuration can be achieved using either two dual-port NICs or four single-port NICs or a combination of the two in each server. Two switches for the cluster interconnect and two hubs or switches for the client LAN connections are shown here. This configuration is the recommended architecture for optimal fault.

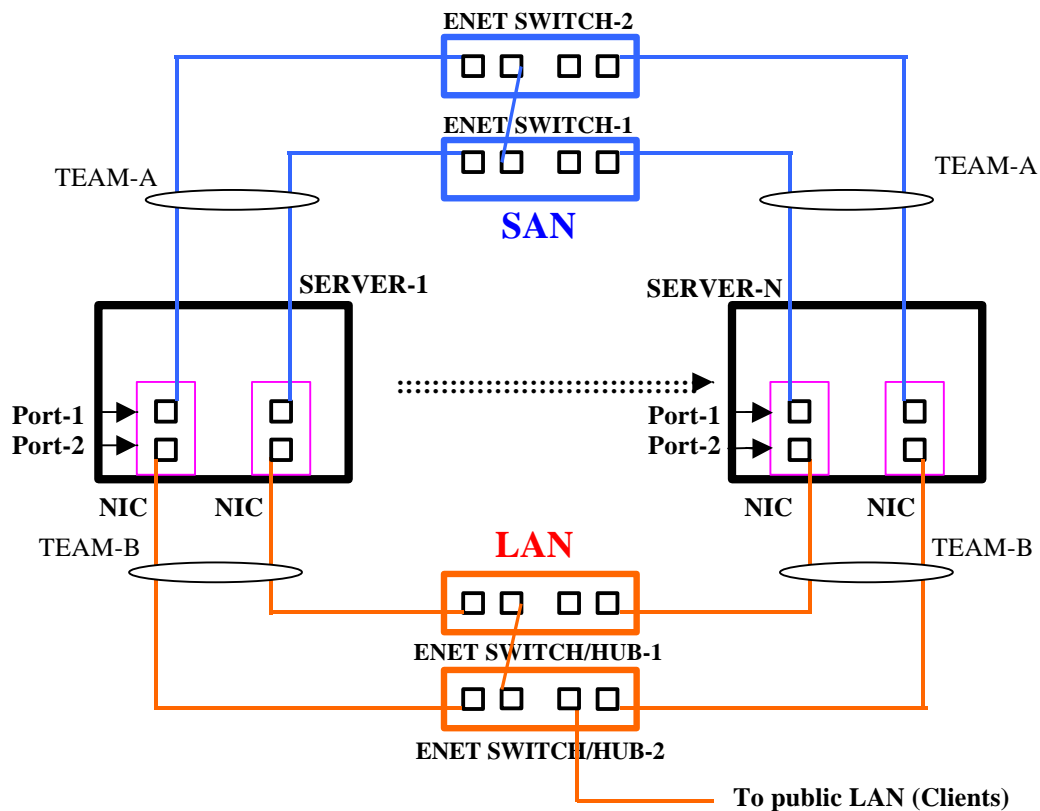


Figure 2. Fault Tolerant Configuration

Implementing this configuration in the NT environment is achieved through the use of an NT networking feature known as “port teaming”. Port teaming is defined as combining or merging of ports to create a standby or active alternate communication path. Follow the instructions below to create a fully redundant and highly available communications network:

Step 1. Connect the top port (Port-1) of one NIC in each server to the first cluster interconnect switch (Enet Switch-1).

Step 2. Connect the top port (Port-1) of the other NIC in each server to the second cluster interconnect switch (Enet Switch-2).

Step 3. Connect the bottom port (Port-2) of one NIC in each server to the first public client LAN switch (Enet Switch/Hub-1).

Step 4. Connect the bottom port (Port-2) of the other NIC in each server to the second public client LAN switch (Enet Switch/Hub-2).

Refer to the “**Steps for configuring NICs utilizing teaming**” portion of this document for the step-by-step teaming instructions.

- Assign the top ports of both NICs in each server to one team (“TEAM-A” see Figure 2.)
- Assign the bottom ports of both NICs in each server to one team (“TEAM-B” see Figure 2.)

NOTE: For the teaming option, the Adaptive Load Balancing (ALB) is strongly recommend — it increases the transmission throughput by supporting up to four NIC teams, with two or four NICs per team.

Teaming Configuration for Higher Availability

Hardware and Software Requirements

- **Hardware**

To create a fault tolerant configuration of the cluster, the redundancy option of Compaq's Fast Ethernet dual-redundant NC series NIC cards is required. To use this function, two dual-port or four single-port (or a combination of both) NICs are required for each server in the cluster.

For the cluster interconnect connection, two network switches are required. Two hubs may be used in a two node configuration. For better performance, the hubs should be replaced by switches when a third node is added.

- **Software**

Oracle Software

Contact Oracle Corp. for all Oracle software.

Compaq Software

Contact Compaq Corp. for Compaq Parallel Database Cluster software and the latest drivers of the Compaq Fast Ethernet NC series NIC cards.

Microsoft SP4 for Windows NT 4.0

To implement the dual redundant option of the Fast Ethernet NIC cards, Microsoft SP4 for Windows NT 4.0 is highly recommended. Install SP4 prior to utilizing the NIC teaming option. Contact Microsoft for this software.

After you team your NICs, you must ensure that the binding order of the ports are correct. This is explained later in the guide.

Note: If you are required to use SP3 (which is not recommended), please contact Lino Costantino @ Compaq Corp. at (408) 285-6983 or lino.costantino@compaq.com.

Teaming Configuration of Ethernet Interconnects for Compaq Fast Ethernet NC Series NIC cards

To achieve higher availability in a cluster, the enhanced dual redundant features of Compaq Fast Ethernet NICs are required.

The dual redundancy option of the Fast Ethernet NICs are achieved by assigning individual NICs/ports to a team. Microsoft Windows NT 4.0 Service Pack 4 (SP4) is required for implementing NIC teaming properly. This “teaming” configures the ports as primary and secondary (or backup) ports with a common IP address. After building a team, identify the type of team (NFT or ALB).

For the best teaming option, the Adaptive Load Balancing (ALB) is strongly recommended — ALB increases transmission throughput by supporting up to four NIC teams, with two or four NICs per team. NICs configured for ALB also provide the benefit of NIC fault tolerance (NFT).

Option 1. NIC Fault Tolerance (NFT) Configuration

NIC Fault Tolerance provides the safety of an additional backup link between the server and hub or switch. In the event of hub or switch port, cable, or NIC failure, you can maintain uninterrupted network performance.

NIC fault tolerance is implemented with a primary NIC and a backup, or secondary NIC. During normal operation, the backup NIC will have its transmission disabled. If the link to the primary NIC fails, the link to the secondary NIC automatically takes over.

- **Setting Up NIC Fault Tolerance**

To use NIC Fault Tolerance (NFT), you must have two or more Compaq Fast Ethernet NC series NICs installed in your server. The default settings of the team properties are recommended.

Option 2. Adaptive Load Balancing (ALB) Configuration

The ALB (Adaptive Load Balancing) option is recommended for the best performance and availability of your cluster.

Adaptive Load Balancing (ALB) is a simple and efficient way to increase your server's transmission throughput. With ALB, as you add NICs to your server, you can group them in teams to provide up to 400 mbps with a maximum of four Compaq Fast Ethernet NICs. The ALB software continuously analyzes transmission loading on each NIC and balances loading as needed.

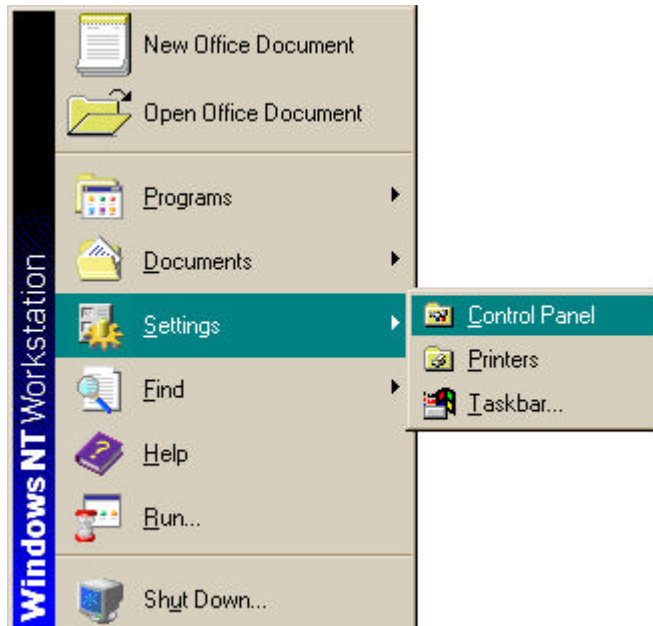
- **Setting Up Adaptive Load Balancing**

To use Adaptive Load Balancing (ALB), you must have two, three, or four Compaq Fast Ethernet NC series NICs installed in your server, all linked to the same network switch. To configure ALB, all the NICs in the team must operate at 10/100 Mb/s or 1 Gb/s. The default settings of the team properties are recommended.

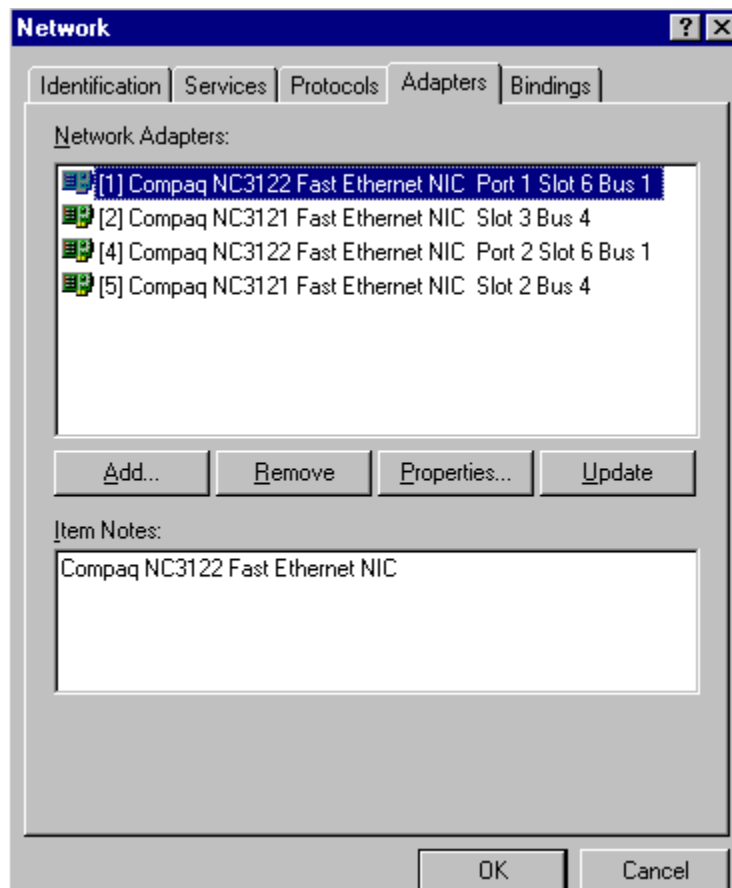
Note: For maximum benefit, ALB should not be used under NetBEUI and some IPX environments. For a list of specific environments, see the your Administrator Guides.

Steps for Configuring NICs Utilizing Teaming

1. Click the **Start** menu and go to **Control Panel** under the **Settings** menu.

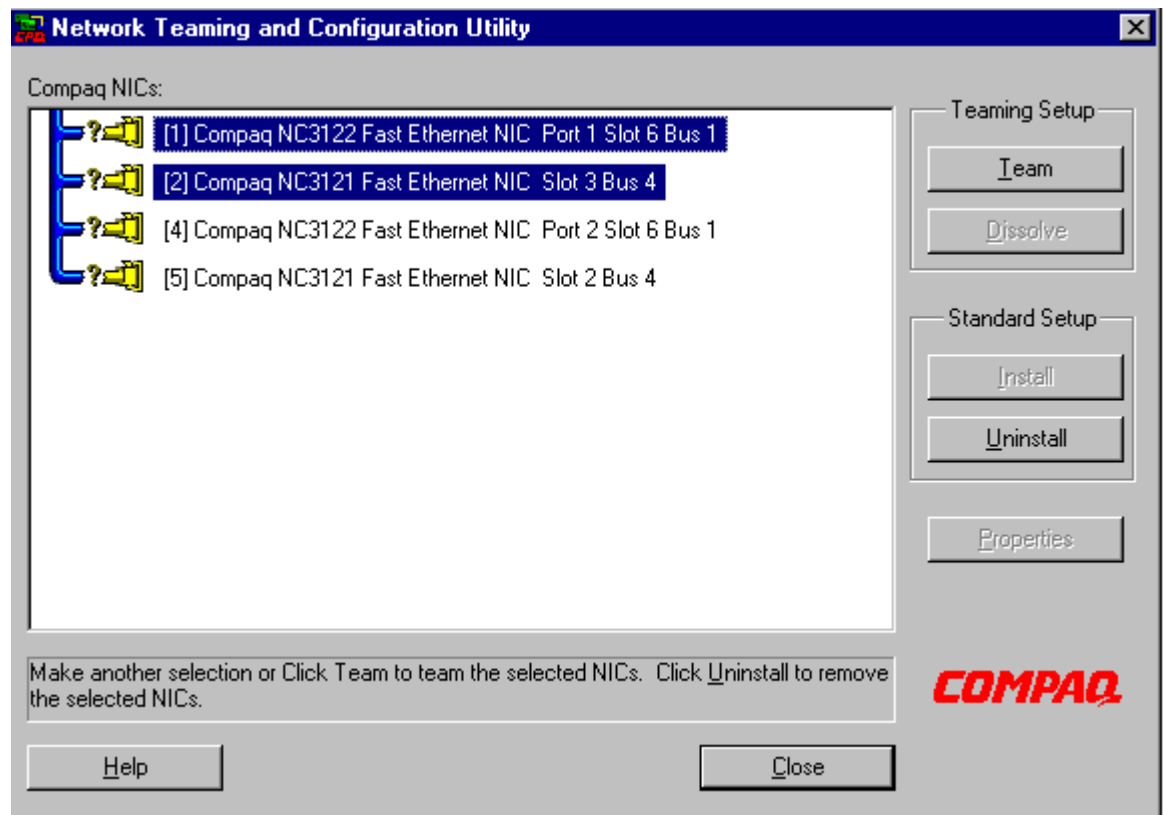


2. Double click the **Network** control panel. Click on the **Adapters** tab. All of the installed NICs should be listed. Select a NIC and click on **Properties**.



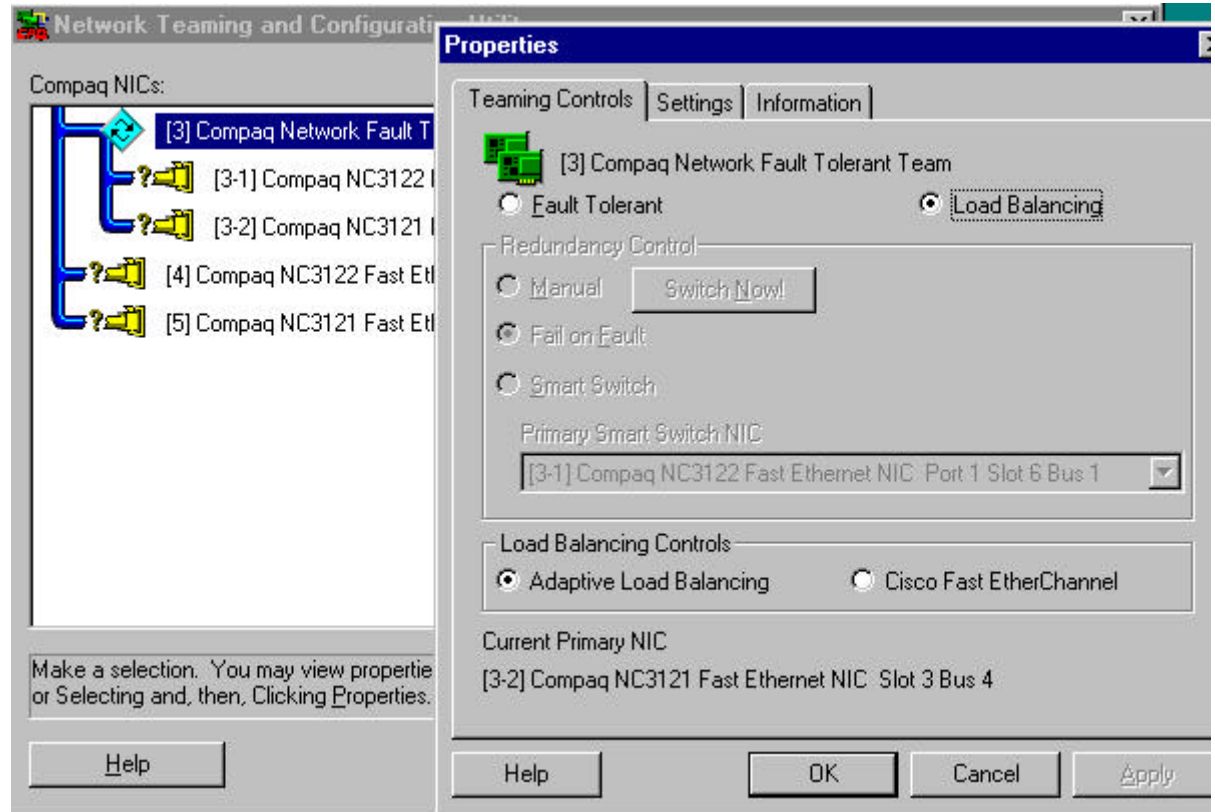
- This takes you to the Compaq **Network Teaming and Configuration Utility** main window. All of the installed NICs are listed again. You are now ready to build a team. To do so, make sure the different NICs are all on the same network and running at the same speed.

To build a new team, highlight the NICs you want in the team and select **Team** under **Teaming Setup**.

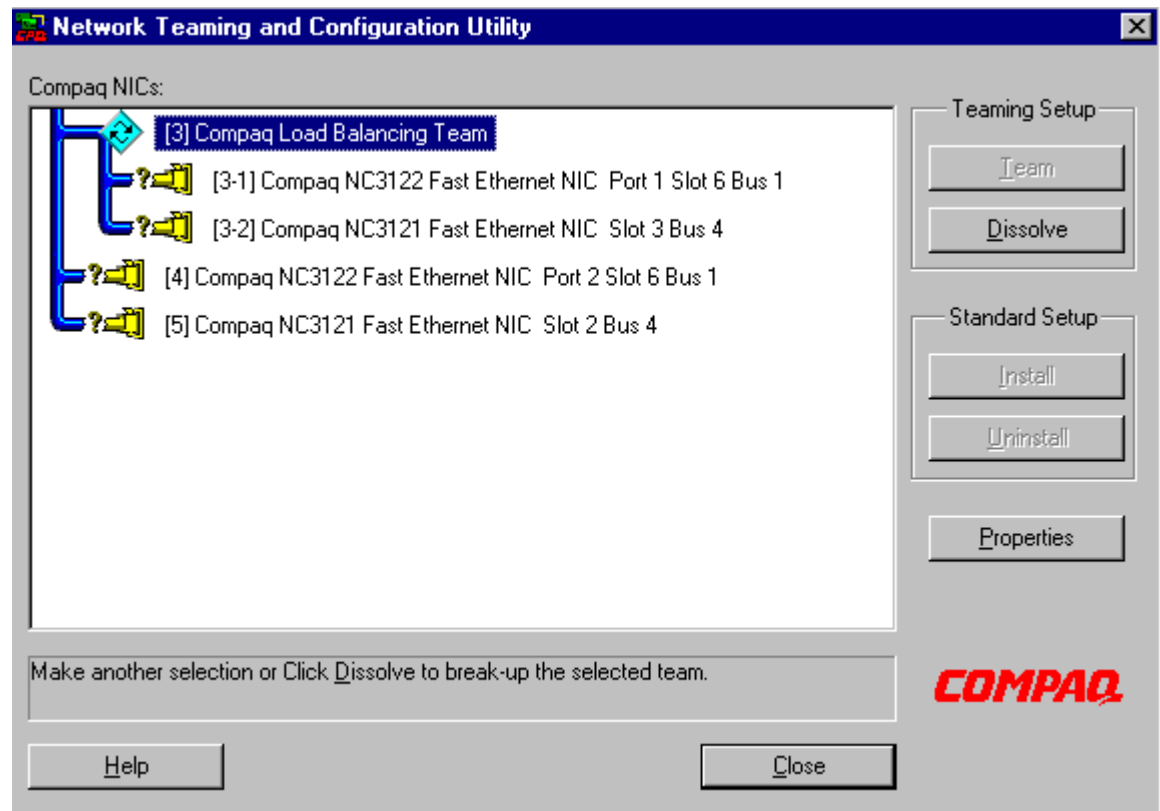


Note: You may **not** add NICs of different speeds to a team that is already established.

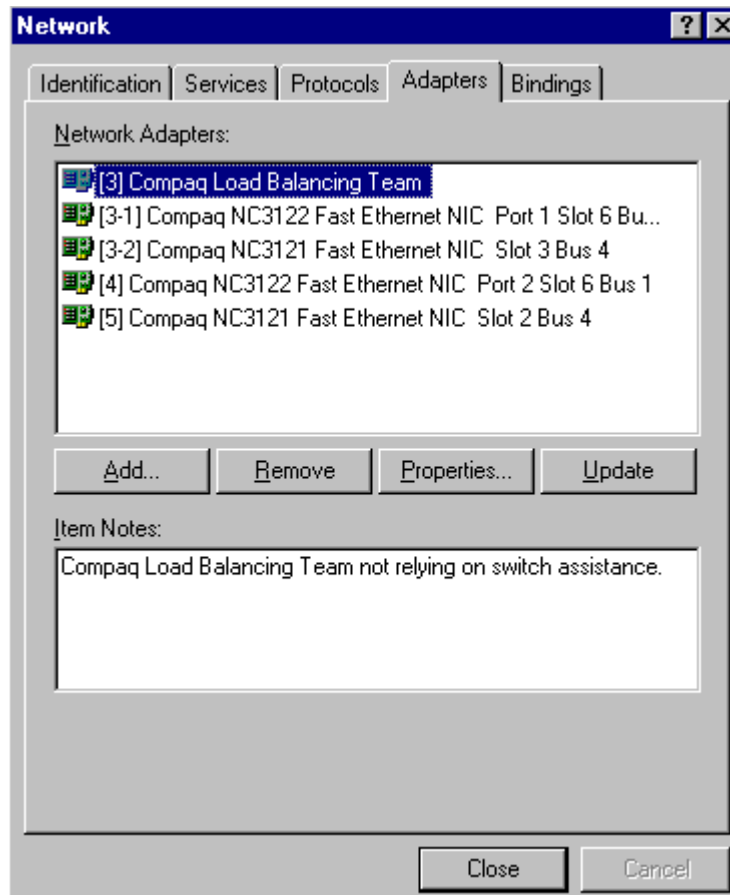
- This takes you to the **Properties** window. Select the **Teaming Controls** tab (default tab). Then select either **Fault Tolerant** for the NFT option or **Load Balancing** for the ALB option. ALB is the recommended option. Click **OK** to finish.



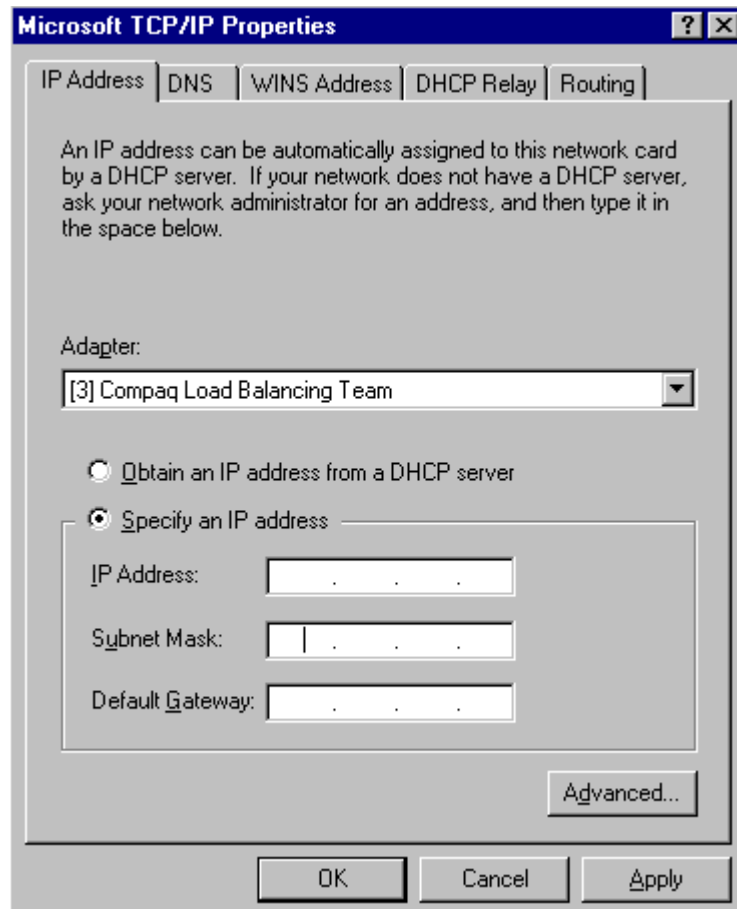
5. This takes you back to the **Network Teaming and Configuration Utility** main window. Click **Close** to complete this operation.



6. This takes you back to the **Adapters** tab in the **Network** control panel. Click **Close**.



7. This takes you to the **Microsoft TCP/IP Properties** window. Select the team that you just created under the **Adapter:** pull down menu and input the proper IP addresses.



Repeat steps to create more teams.

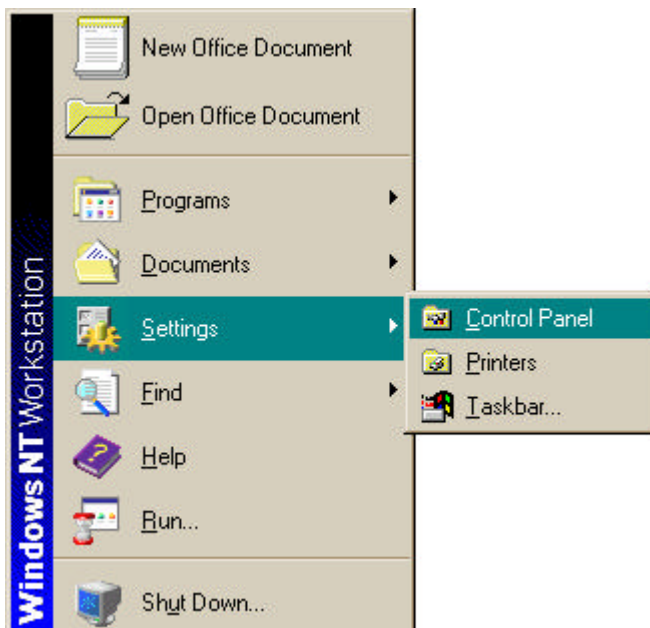
8. When prompted, restart your computer.

Verifying Client LAN and Cluster Interconnect IP Addresses

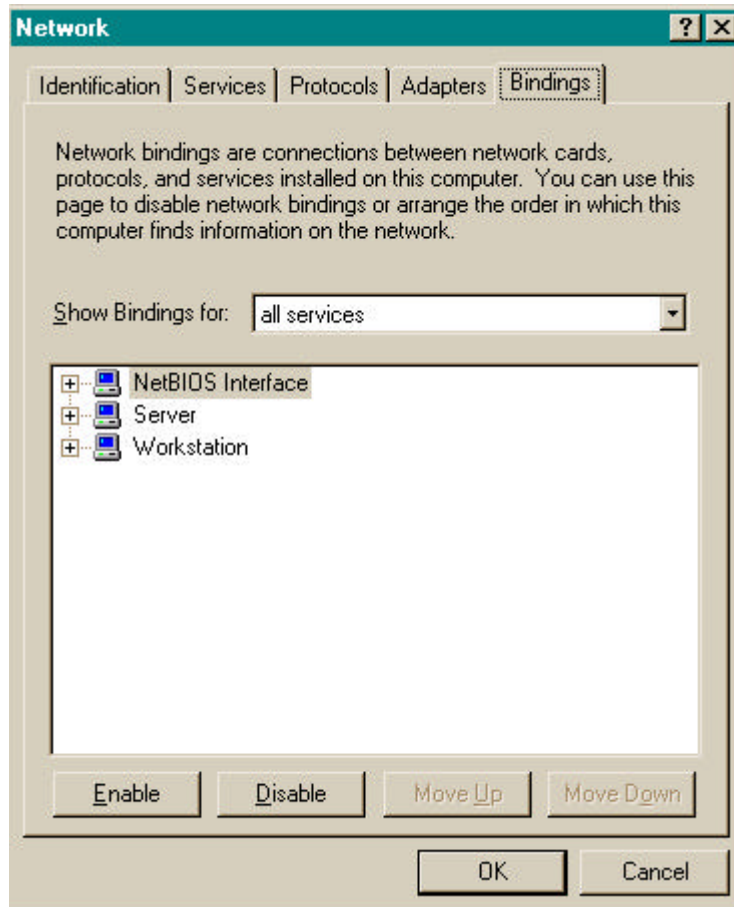
Verify that the IP addresses for the client LAN and cluster interconnect are correctly assigned by pinging the machine host name. (You can find the machine host name by selecting the **Identification** tab in the **Network** control panel.)

The IP address returned when you ping the machine host name should be the client LAN IP address. If the IP address returned by the ping command is not the client LAN IP address you specified, perform the following steps to correct binding order of the NICs.

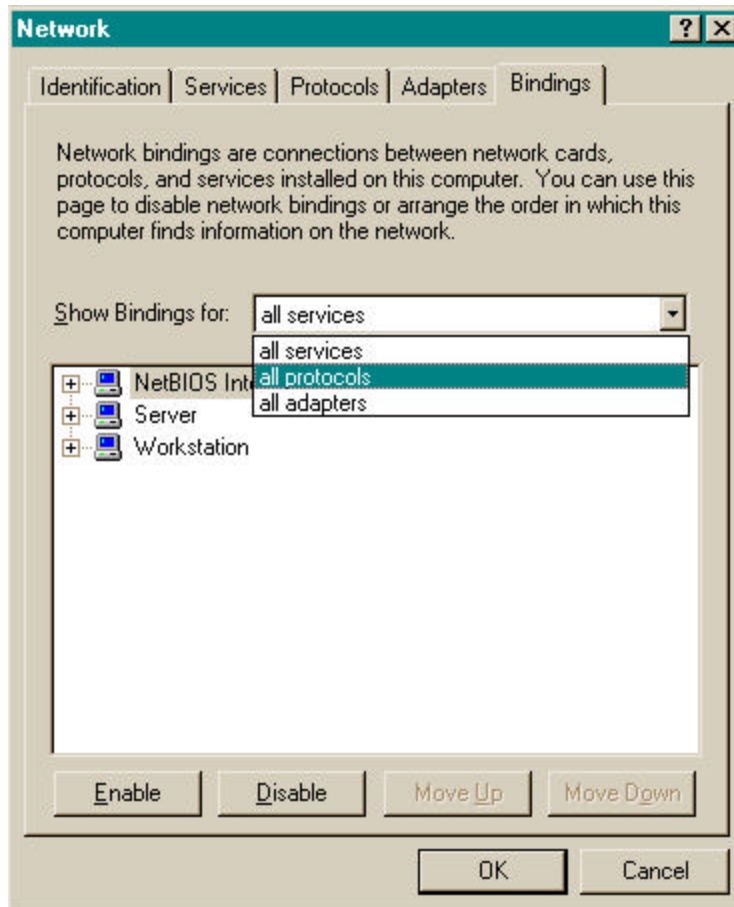
- 1) Click the **Start** menu and go to **Control Panel** under the **Settings** menu.



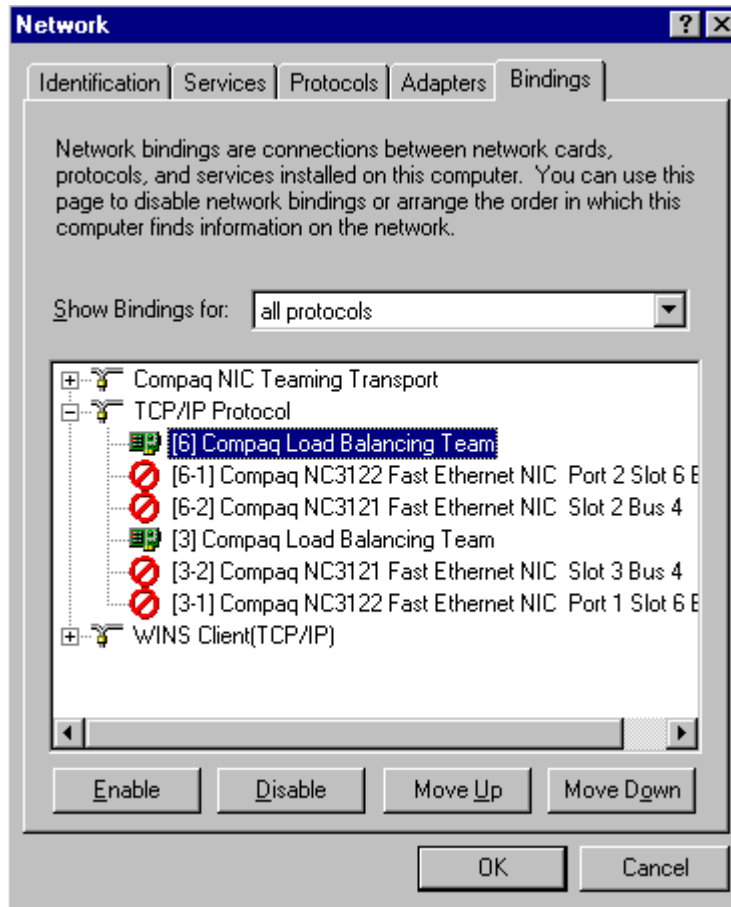
2) Double click the **Network** control panel. Select the **Bindings** tab.



- 3) Select **all protocols** from “**Show Bindings for:**” pull down menu.



- 4) Click + (plus sign) next to **TCP/IP Protocol**. A list of all installed NICs/teams appear.



Windows NT Server uses the IP address specified for the NIC/team at the **top** of the list in this window as the client LAN IP address. If the NIC/team that you assigned as the client LAN (by specifying the IP address in step 7 of previous section of this paper) is not at the top of the list, find the NIC/team you specified for the client LAN in the list and select it. Click the **Move Up** button to position this NIC/team at the top of the list.

- 5) Click **OK** to finish operation. When prompted, restart your computer.