

February 2002
16D2-0202A-WWEN (nt)

Prepared by Industry Standard Server
Group (Cupertino High Availability
Integration and Test Division)

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
Configuring NIC Teaming for Compaq's Fast or Gigabit Ethernet NC Series Adapters ..	9
Steps for Configuring NIC Teaming	11

Supported Ethernet Interconnects for Compaq Parallel Database Clusters using Real Application Clusters

Abstract:

This Technical Guide reviews the supported Ethernet client LAN and cluster interconnect options for Compaq Parallel Database Clusters with Real Application Clusters (RAC). Hardware and software requirements are discussed along with the limitations inherent to them. Both redundant and non-redundant configurations are explained.

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.

OPS, and RAC are registered trademarks of Oracle Corp.

Copyright ©2002 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

First Edition (February 2002)

Document Number 16D2-0202A-WWEN (nt)

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 (PDC) with Real Application Clusters (RAC). This Guide reviews the supported Ethernet architectures for the client LAN and cluster interconnect (sometimes known as System Area Network or SAN) in a RAC environment. It does not provide detailed cluster network configuration information. Please refer to the Order and Configuration Guide and the PDC 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 architectures for both the client LAN and cluster interconnect. Limitations and requirements are outlined.

Two options are detailed for the implementation of fully redundant client LAN and cluster interconnect connections: Network Fault Tolerance (NFT) and Transmit Load Balancing (TLB), formerly known as Adaptive Load Balancing (ALB). TLB is recommended for the best performance and availability of your cluster.¹ Hardware and software configuration requirements will be defined for these architectures. The stated configurations below are supported for Compaq Parallel Database Clusters in a RAC environment.

¹ Transmit Load Balancing (TLB) is only supported with either 10 or 100MB Ethernet and NOT 1000MB Ethernet for the cluster interconnect. If you are using TLB for the client LAN connections then 10, 100 or 1000 Mbps is supported.

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 five hardware failures are primary points in which a disruption of communication can occur:

- A port on a single or 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 for both the cluster interconnect and the client LAN connections in a fully redundant configuration.

The cluster interconnect remains available if:

- A network controller port used for the primary cluster interconnect path fails.
- A network controller used for the primary cluster interconnect path fails.
- The network cable(s) used for the primary cluster interconnect path fails.
- The port(s) on the switch used for the primary cluster interconnect path fails.
- The switch used for the primary cluster interconnect path fails.

Single 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 client LAN path fails.

Single 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 single point of failure in the communications network. Certain non-fault tolerant Ethernet architectures are supported.

For two or more node configurations, a switch is required but it is not necessary to be a fully redundant configuration in order to be supported (See Figure 1).

NOTE: As the number of nodes increase in a clustered configuration, the higher the chances of a single point of failure.

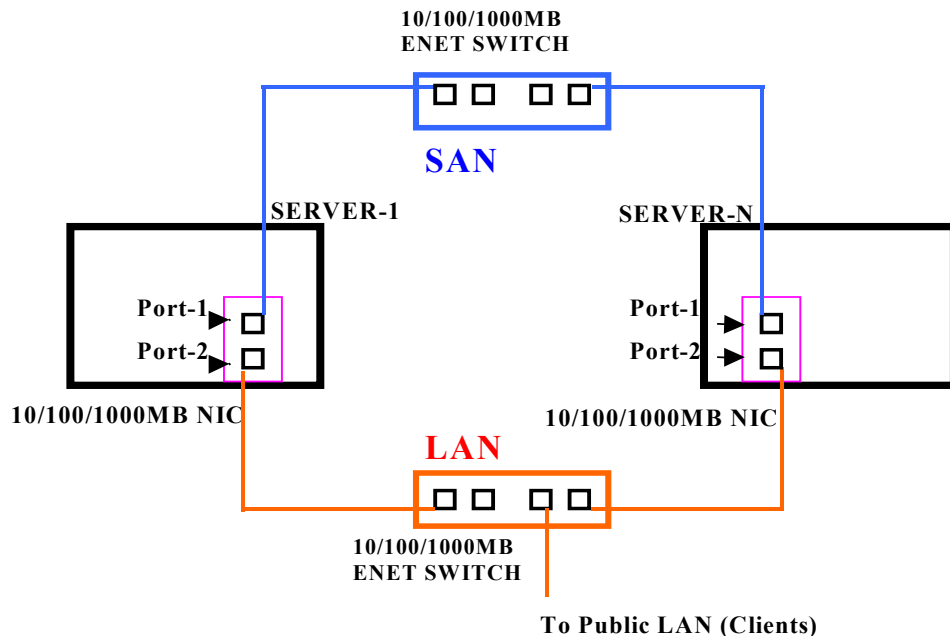


Figure 1: Example of a non-fault tolerant, 2-node configuration using a switch for both client LAN and cluster interconnect.

Supported Fault Tolerant Architecture

The supported fault tolerant architecture combines the use of Microsoft Windows 2000 Advanced Server, Compaq software, and Compaq Fast or Gigabit 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 primary path and a secondary or standby path on separate NICs (either dual or single port), for each node. This path/port teaming feature is enabled with a Compaq Teaming utility in the Microsoft Windows 2000 Advanced Server network control panel.

Several configurations are supported for the cluster interconnect and the client LAN. In this document, only the most robust configuration will be provided for highest availability and performance (See Figure 2).

NOTE: This document will address redundancy in both the client LAN and cluster interconnect. It is not required that both networks be fully redundant to be supported. Any combination will work as well.

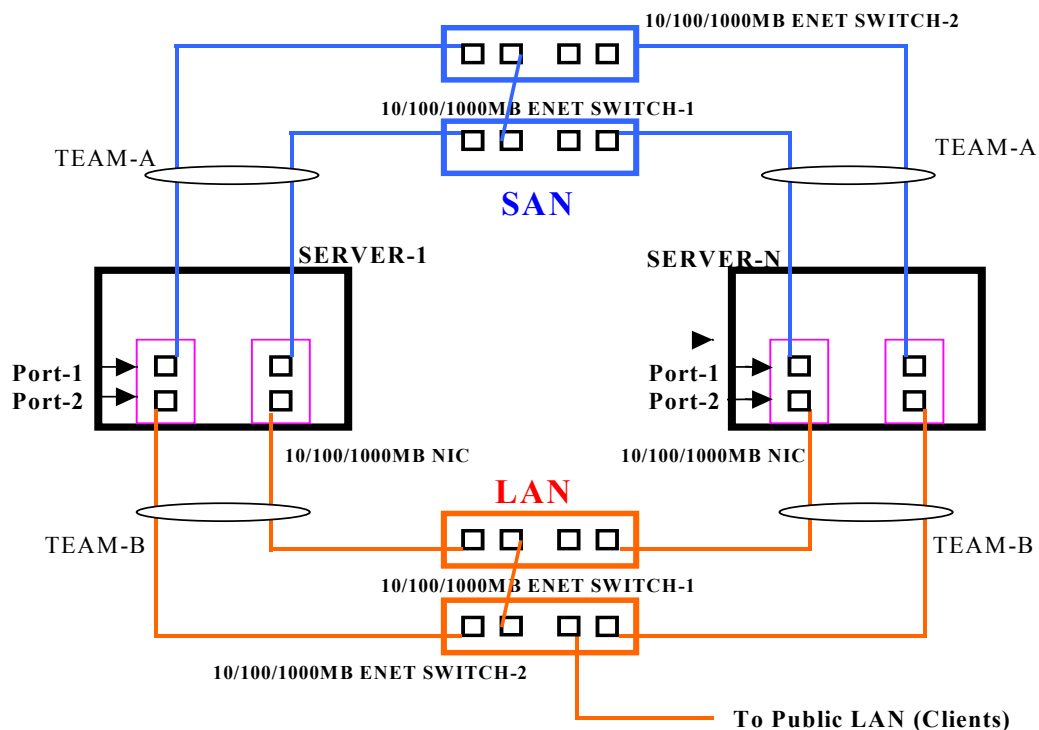


Figure 2: Fault Tolerant Configuration

Implementing this configuration in the Windows 2000 Advanced Server environment is achieved through the use of “port teaming.” Port teaming is defined as the combining or merging of adapter ports to create a single “virtual” port. Through Compaq’s Network and Teaming Configuration Utility, this virtual port can be set to the Fault Tolerant (NFT) or Load Balance (TLB) option.

Step 1. Connect the top port (Port-1) of one NIC in each server to the first cluster interconnect switch (10/100/1000MB Enet Switch-1).

Step 2. Connect the top port (Port-1) of the other NIC in each server to the second cluster interconnect switch (10/100/1000MB Enet Switch-2).

Step 3. Connect the bottom port (Port-2) of one NIC in each server to the first public client LAN switch (10/100/1000MB Enet Switch-1).

Step 4. Connect the bottom port (Port-2) of the other NIC in each server to the second public client LAN switch (10/100/1000MB Enet Switch-2).

NOTE: Refer to the “**Steps for Configuring NIC Teaming**” portion of this document for the walkthrough on using Compaq’s Network and Teaming Configuration Utility.

Step 5. Assign the top ports of both NICs in each server to one team (“TEAM-A” see Figure 2).

Step 6. Assign the bottom ports of both NICs in each server to one team (“TEAM-B,” see Figure 2).

NOTE: For teaming, Transmit Load Balancing (TLB) is strongly recommended.² TLB increases the transmission throughput by supporting up to eight NIC teams with two to eight adapter ports per team.

² Transmit Load Balancing (TLB) is only supported with either 10 or 100MB Ethernet and NOT 1000MB Ethernet for the cluster interconnect. If you are using TLB for the client LAN connections then 10, 100 or 1000 Mbps is supported.

Teaming Configuration for Higher Availability

Hardware and Software Requirements

- **Hardware:**

To create a fault tolerant configuration of the cluster, Compaq's Fast Ethernet or Gigabit Ethernet NC series NICs are required. For proper redundancy configuration, two dual-port or four single-port (or combination of both) NICs are required for each server in the cluster.

For the cluster interconnect and client LAN connections, two network switches are required.

- **Software:**

Oracle Software

Contact Oracle Corporation for all Oracle Software.

Compaq Software

Contact Compaq Computer Corporation for Compaq Parallel Database Cluster software. The latest drivers for Compaq's Fast and Gigabit Ethernet NC series NIC cards, and the latest drivers for Compaq Teaming are downloadable from the Compaq Website at the following location: <http://www.compaq.com/support/files/networking/nics/index.html>. For NIC teaming to work properly use SmartStart 5.20 or later.

Microsoft Software

To implement NIC teaming properly with either Fast or Gigabit Ethernet NIC cards, one must install Windows 2000 Advanced Server with Service Pack 2. Make sure that Service Pack 2 has been installed prior to utilizing the NIC teaming option. Download SP2 from this location: <http://www.microsoft.com/windows2000/downloads/servicepacks/default.asp>.

NOTE: After implementing NIC teaming, ensure that the binding order of the ports are correct. This is explained later in this guide.

Configuring NIC Teaming for Compaq's Fast or Gigabit Ethernet NC Series Adapters

In order to achieve higher availability in a cluster, Compaq's Fast or Gigabit Ethernet NICs are required.

Redundancy is achieved by merging ports from Compaq's Fast or Gigabit Ethernet NICs to form a team. Microsoft Windows 2000 Advanced Server Service Pack 2 (SP2) is required for implementing NIC teaming properly. Teaming combines chosen adapter ports and designates a primary and a secondary or standby port. The merged ports are then assigned a common "virtual" IP address. After deciding on which ports will be teamed, identify the type of teaming that will be used (NFT or TLB).

Option 1. Network Fault Tolerance (NFT) Configuration

NFT provides the safety of an additional backup link between the server and switch. In the event of switch, switch port, cable, NIC port or NIC failure, uninterrupted network performance can be maintained. NFT supports two to eight network adapter ports per team.

NFT is implemented with a primary NIC port and a standby NIC port. During normal operation, the standby NIC port will have its transmission disabled. If the link to the primary NIC port fails, the link to the standby NIC port automatically takes over.

- **Configuration of Network Fault Tolerance For An Oracle 9i RAC Environment**

To use NFT, install two or more Compaq Fast or Gigabit Ethernet NC series NICs in the server. If using the Fault Tolerant teaming option for the cluster interconnect, all NICs in the team must operate at either 10, 100 or 1000 Mbps transfer rates. If using the Fault Tolerant teaming option for the client LAN connections, a combination of 10, 100 and 1000 Mbps transfer rates across the NICs in a team can be used. The default settings of the team properties are recommended.

If "Supported Fault Tolerant Architecture" is used, make sure that spanning tree is **enabled** on the ports that connect the switches (whichever switch ports the crossover cable is using). However, spanning tree should be **disabled** on all switch ports to which a Compaq Network Adapter Team is attached.

Option 2. Transmit Load Balancing (TLB) Configuration

The TLB option is recommended for the best performance and availability of your cluster.³

TLB is a simple and efficient way to increase the server's transmission throughput. TLB supports two to eight network adapter ports per team. With TLB, a maximum of eight Compaq Fast or Gigabit Ethernet NICs can be added to a team, providing up to a maximum theoretical transmit throughput of 800 Mbps for Fast Ethernet or 8000 Mbps for Gigabit Ethernet.

³ Transmit Load Balancing (TLB) is only supported with either 10 or 100MB Ethernet and NOT 1000MB Ethernet for the cluster interconnect. If you are using TLB for the client LAN connections then 10, 100 or 1000 Mbps is supported.

- **Configuration of Transmit Load Balancing For An Oracle 9i RAC Environment**

To use TLB, two or more Compaq Fast or Gigabit Ethernet NC series NICs must be installed in the server. If using the Load Balancing teaming option for the cluster interconnect, all the NICs in the team must operate at either 10 or 100 Mbps. If using the Load Balancing teaming option for the client LAN connections, all the NICs in a team can operate at either 10, 100 or 1000 Mbps. For the cluster interconnect and client LAN connections make sure to choose **Balance with IP Address** under **Load Balancing Controls**. The default settings of the team properties are recommended.

If “Supported Fault Tolerant Architecture” is used, make sure that spanning tree is **enabled** on the ports that connect the switches (whichever switch ports the crossover cable is using). However, spanning tree should be **disabled** on all switch ports to which a Compaq Network Adapter Team is attached.

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

Steps for Configuring NIC Teaming

1. Click the **Start** menu and go to **Control Panel** under the **Settings** menu.
2. Double click on **Compaq Network**.
3. The **Compaq Network Teaming and Configuration Properties** screen is displayed. The list shows all the installed NICs on the server. For teaming to work properly, make sure the NICs being teamed are all on the same network and running at the same transfer rate.

To build a new team, highlight the desired NICs for the team and click **Team** under **Teaming Setup**.

4. The Properties dialog box will appear.

Select the **Teaming Controls** tab (default tab). There are two choices:

- **Fault Tolerant** for Network Fault Tolerance.

If choosing **Fault Tolerant**, choose between **Fail on Fault**⁴ (default) and **Smart Switch**⁵.

- **Load Balancing** for Transmit Load Balancing.

If choosing **Load Balancing**, choose **Adaptive Load Balancing** (default) and **Balance with IP Address** under **Load Balancing Controls**.

TLB is the recommended option.⁶

Click **OK** to finish.

NOTE: The primary adapter in a NFT team is shown as the first listed adapter and shows a complete blue bar rather than a shaded blue bar. The Primary adapter in a TLB team is the first listed adapter in the Load Balanced team.

5. When at the **Compaq Network and Configuration Properties** screen, click the **OK** button.
6. One or more **Digital Signature Not Found** messages will appear. Click **Yes** to all of them to continue.
7. An information dialog box will appear reminding you to make sure that all Protocols and Services are properly configured for the Network Team. Click **OK**.

⁴ System will switch over when a fault is detected on the active NIC. The standby NIC becomes active and remains active.

⁵ You select the preferred primary NIC of the team. This is always the active NIC as long as it is operational. Should the preferred primary NIC fail, it will fail to the standby NIC. However, once the preferred primary NIC is repaired, the Teaming Utility will automatically restore it as the active NIC.

⁶ Transmit Load Balancing (TLB) is only supported with either 10 or 100MB Ethernet and NOT 1000MB Ethernet for the cluster interconnect. If you are using TLB for the client LAN connections then 10, 100 or 1000 Mbps is supported.

8. Restart the server for the changes to take effect. Click **Yes** to reboot the server.
9. Once the server restarts, Click the **Start** menu and go to **Control Panel** under the **Settings** menu.
10. Double click **Network and Dial-up Connections**. Make sure the “virtual” IP icon shows up with respect to the team(s) just created.

Double click the “virtual” IP icon and click on **Properties**. Scroll down and select **Internet Protocol (TCP/IP)** then click on **Properties**. Input the proper IP address and subnet mask. Repeat steps with all teams created.

NOTE: Make sure the bindings are set properly by going to the **Advanced** tab and selecting **Advanced Settings**. Under the **Adapters and Bindings** (default) tab, make sure the LAN connection is first on the list followed by the cluster interconnect. To test that the settings have been saved, ping the local host by hostname. The IP address set for the LAN connection should be returned.

11. Open the **Event Viewer**. Clear the **System Log**.
12. Remove the network cable from the back of the primary NIC within either the NFT or TLB Team just created. An error message that the cable is unplugged will be generated. Refresh the Event Viewer and note the message. Through the **Compaq Network Teaming and Configuration Properties** screen, verify that the unplugged cable shows up as a disconnected cable icon.
13. Plug the cable back into the NIC. Through the **Compaq Network Teaming and Configuration Properties** screen, verify the disconnected cable icon changes showing that the cable is now reinserted. Refresh the **Event Viewer** and note the messages.