

WHITE PAPER

June 1996

Prepared By
OS Development

Compaq Computer
Corporation

CONTENTS


| | |
|--|---|
| Compaq's Network Device Drivers | 3 |
| Hardware Failure Events | 3 |
| Network Status Change Events | 4 |
| Driver Recovery Process | 5 |
| Configuring NetFlex-3 with SCO OpenServer 5.0 | 6 |
| Advanced Network Control Utility | 6 |
| Configuration Tips | 8 |
| Redundant Controller Pair Hardware Configuration | 8 |
| Tunables Configuration | 8 |

Advanced Network Fault Detection and Correction Feature for UNIX

Compaq has a twofold solution to ensure high network availability of your Compaq server. The first line of defense is Compaq's network device drivers, which have built-in recovery mechanisms. These mechanisms allow recovery from many hardware failure events, network status events and other network errors. When the device driver detects a network controller failure, it reports the failure and attempts to restart the failed controller.

The second part of the solution is the Compaq Advanced Network Control Utility. This utility allows you to configure two network controllers (a primary and a secondary) as a Redundant Controller Pair. Initially, the primary is the active controller and the secondary is the standby controller. The standby controller does not handle any network activity unless a failure is detected on the active controller. If a failure is detected on the active controller, the driver switches to the standby controller.

This paper describes the recovery features provided by the Compaq network drivers as well as configuration recommendations and diagnostic processes to help you recover from network availability problems.



NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL.

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. Compaq does not warrant products other than its own strictly as stated in Compaq product warranties.

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

Compaq and NetFlex, registered United States Patent and Trademark Office.

Netelligent is a trademark and/or service mark of Compaq Computer Corporation.

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

©1996 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

High Availability Features for Compaq NetFlex-3 Controllers

First Edition (June 1996)

386A/0696

COMPAQ'S NETWORK DEVICE DRIVERS

All Compaq network device drivers provide integrated error recovery that allow the device drivers to detect failure events and recover from these errors. The drivers can reset the network interface card and continue running, usually without noticeable disturbance to normal operation. The drivers can recover from the following error types:

- Hardware Failure Events
 - Adapter Check Interrupt
- Network Status Change Events
 - Link Status Change
 - Transmit Integrity Check Failure

Hardware Failure Events

Hardware failures can occur, and it is important to recognize the factors that indicate a problem. The cet driver attempts recovery based on certain events described below.

Adapter Check Interrupt

Adapter Check errors indicate that the hardware detected a problem and reported it to the driver. A console error message that details the problem is generated and an immediate attempt to recover is begun. The Adapter Check Interrupts statistic, viewable under the Advanced Network Control Utility (cpqnm) Diagnostics menu selection, will be incremented.

Driver Recovery Process

The NetFlex-3 and Netelligent controller drivers implement a recovery process that may be initiated whenever fatal network status changes occur or by manual user intervention. These two cases are known as Automatic Recovery Mode and Manual Recovery Mode, respectively. You may switch between Automatic and Manual modes using the Operations menu item in the cpqnm control program.

Automatic Recovery Mode

When a driver interface is configured for Automatic Recovery Mode, the driver acts automatically to recover from any fatal errors it detects on that interface. No intervention is required and the driver will continue to attempt recovery until it is successful or until the process is terminated (driver shutdown). Automatic Recovery Mode is the default for all driver interfaces.

If an interface experiencing a failure is not configured in a Redundant Controller Pair, the driver will reset the interface. Following the interface reset, the driver re-enters its fault-detection process for the interface and continues to attempt recovery if the failure persists.

If an interface experiencing a failure is the active interface in a Redundant Controller Pair, the driver switches the roles of the active and standby interfaces (standby becomes active). If this attempt is successful, the driver will continue to operate in this configuration until another fault is detected on the current active interface. If the attempt to switch control to the standby is not successful, the driver will toggle back and forth between the pair until one of them successfully restarts. Once one of the interfaces responds, the driver will begin the recovery process for the current standby interface. Recovery of the standby involves resetting the interface until it is fully functional.

Manual Recovery Mode

When a driver interface is configured for Manual Recovery Mode, the driver makes no attempt to recover from faults of any kind. This mode is intended for users who either wish to disable fault recovery features of the driver or test the fault recovery features manually. Depending on whether or not an interface is configured in a Redundant Controller Pair, you may manually reset an interface or force a switch to a standby interface.

If the Redundant Controller Pair is in Manual Recovery Mode and the standby controller is highlighted, the only option available is the restart option. This option may be used to test the interface reset or to perform other troubleshooting tasks.

If the Redundant Controller Pair is in Manual Recovery Mode and the active controller is highlighted, both the reset option and the ability to force a switch to the standby interface are available. These options may be used to test the operation of either controller or to perform other troubleshooting tasks.

CAUTION: If an interface is set to Manual Recovery Mode via the cpqnm application and you exit cpqnm without restoring the interface to Automatic Recovery Mode, all fault-detection and recovery features will be disabled for that interface. Before exiting cpqnm, be certain that each interface is left in the intended recovery mode.

CONFIGURING NETFLEX-3 WITH SCO OPENSERVER 5.0

Advanced Network Control Utility

Compaq Advanced Network Control features are provided by the NetFlex Interface Manager application (cpqnim). Use this utility to configure and monitor the status of the supported controllers. When using cpqnim keep in mind that when changes are made they take effect immediately.

The application displays the current configuration for each of the supported controllers it detects in the system. If the interface is configured in the UNIX kernel, its name (cet[0-7]) will appear at the far left of the display line. If no interface name appears, the interface is not currently controlled by the Compaq driver (unless it is the secondary interface in a Redundant Controller Pair).

The application displays the current status for each of the supported controllers it detects in the system. Since the application operates in both character and graphical (Xwindows) modes, status is indicated differently for each mode. In graphical mode, the interface status is represented by a line with colored dots at each end. Green dots at each end indicate that the interface is operational and has a stable connection to the network. Red dots at each end indicate that the interface has failed or that its network connection has become unstable. In character mode, these conditions are indicated by '*' characters (operational) or 'x' characters (failed).

If one or more Redundant Controller Pairs are established, this relationship is depicted by erasing the name of the secondary interface, indenting its board icon to the right and connecting the two interfaces by an icon specific to the display mode (graphical or character). The label "Active" appears to the right of the interface that is currently active in the pair.

Main Menu Options

The menu options provide you with access to Configuration, Operations and Diagnostics information specific to the interface currently selected. Selecting any of the menu options results in the display of a pulldown menu containing pushbuttons for each of the available functions. If a pushbutton is highlighted, the function it represents is currently available for the selected interface.

Configuration Option

The configuration option allows you to view or modify aspects of a controller's configuration, including binding/unbinding Redundant Controller Pairs and modifying tuneable parameters.

The **<Bind Redundant Pair>** pushbutton allows you to bind a secondary interface to the selected interface as its backup partner. This function will be available only if the selected interface is operational and there are one or more configured interfaces that are eligible to become secondary interfaces (they are configured and have no protocol stacks connected to them). If the function is available, a list of the possible partner choices will be displayed when the **<Bind>** pushbutton is selected. Select an interface from the list and then select the **<OK>** pushbutton. The display will immediately be updated to indicate that the binding has been completed.

The **<Unbind Redundant Pair>** pushbutton allows you to unbind a Redundant Controller Pair. This function is only available if the selected interface is the primary of a Redundant Controller Pair. If the function is available, select the **<Unbind>** pushbutton to unbind the pair. The display will immediately be updated to indicate that the unbinding has been completed.

The **<Tune Parameters>** pushbutton allows you to adjust some of the parameters that control fault detection. See the **Tunables Configuration Tips** section for more information.

Operations Option

The operations option allows you to perform manual operations on an interface, such as resetting or forcing a switch from the active interface in a Redundant Controller Pair to its standby interface.

The **<Auto>** pushbutton places the interface in Automatic Recovery Mode (the default) when selected. When an interface is in Auto mode, only the **<Manual>** pushbutton will be available.

The **<Manual>** pushbutton places the interface in Manual Recovery Mode, disabling all fault-detection features. Selecting the **<Manual>** pushbutton also causes certain other operations to become available (For example, Restart and Switch).

The **<Restart>** pushbutton causes an interface to be restarted. This is exactly the same operation that would take place if the driver detected a fault while in Automatic Recovery Mode.

The **<Switch>** pushbutton causes the driver to swap the roles of the active and standby interfaces in a Redundant Controller Pair. This function is only be available if the selected interface is the configured primary in a Redundant Controller Pair. Once again, this is exactly the same operation that would take place if the driver detected a fault on the current active interface of a Redundant Controller Pair set for Automatic Recovery Mode.

Diagnostics Option

The diagnostics option allows you to view diagnostic information for the selected interface.

The **<Condition>** pushbutton displays a set of state variables and event counters.

Configuration Tips

Redundant Controller Pair Hardware Configuration

- Supported controllers:
 - NetFlex-3 EISA and PCI Controllers
 - Netelligent 10 T PCI UTP Controllers
 - Netelligent 10/100 TX PCI UTP Controllers
- The secondary controller must be connected to the same network segment (subnet) as the primary controller, but you can connect the secondary controller to a different hub. Should the primary controller failure be related to a hub problem, the secondary controller will continue operation. For example, if two physical networks connected by a MAC bridge comprise a single network segment (subnet), the two controllers may be connected to hubs on either side of the MAC bridge. Note that it does not matter what network topology exists on either side of the bridge (Example: 100VG-AnyLAN on one side, 10Base-T on the other), providing both sides are part of the same logical network segment.
- The Transmitter Integrity Check will not be able to detect link failures on NetFlex-3 controllers with older (Bit-rate) 10Base-T PHY modules when operating in full duplex mode.

Tunables Configuration

The tuneable parameters used in fault detection may require optimization for your particular network. For example, CSMA/CD networks may suffer from “capture effect” when heavily loaded. This means that stations on the network may be prevented from transmitting frames for long periods of time due to other stations’ heavy use of network bandwidth. Since the Transmitter Integrity Check depends on being able to transmit LLC Test frames in order to determine if a network connection is stable, this problem must be addressed.

Tunables have been provided to allow flexibility in dealing with conditions specific to a particular network. These tunables control the number of Test frames used and the time gap between attempts to transmit them. Each tuneable can be increased experimentally until a combination is found that works with a particular network.

Example: An interface connected to a CSMA/CD network is reporting “network connection lost” messages on the system console. It is currently configured to use 2 Test frames with a 5 second delay following each frame. To help reduce or eliminate the error messages, increase the number of Test frames to 3 and increase the delay following each to 10 seconds.

Once again, this may require experimentation to get exactly the right fault detection accuracy for a particular network.