WHITE PAPER

May 1995

Prepared by
Systems Division

Compaq Computer
Corporation

CONTENTS

Compaq Standby Recovery Server

EXECUTIVE OVERVIEW

The Compaq Standby Recovery Server offers minimum downtime for customers with file servers and application servers at remote locations or branch offices without extensive on-site technical expertise. With the Compaq automated switchover process, a second server becomes the active server and is back online in a matter of minutes—without administrator intervention. The Compaq Standby Recovery Server increases system availability for Microsoft Windows NT or NetWare environments by allowing one Compaq ProLiant server to act as a standby for another identically configured ProLiant server.

In the Standby Recovery Server configuration, two ProLiant servers are attached to a common set of ProLiant Storage Systems with a single copy of the operating system, applications, and data. If the primary server fails, each ProLiant Storage System automatically switches from the primary to the recovery server. The recovery server then boots, and the system is back online in minutes without administrator intervention.

The Standby Recovery Server supports Microsoft Windows NT 3.5X, NetWare 3.12, and NetWare 4.X operating systems on Compaq ProLiant Servers, either tower or rack-mounted units.

The Standby Recovery Server offers customers:

- Increased system availability as a result of automated, fast server recovery
- Fully automated server switchover, ideal for unattended, or “lights out,” operation
- Ability to schedule service on a failed server at the most convenient time
- Affordability, even for implementation at numerous operating locations
- Tower or rack-mounted server configurations

INTRODUCTION

As information systems are increasingly used to make real-time decisions, every minute of system downtime becomes more expensive. To reduce downtime in the event that a server fails, some customers use only external storage and they purchase a spare server. With a spare server already in place, the system administrator or a service technician can disconnect the external storage from a failed server, reconnect it to the spare server, and then boot the spare server to restore service.

This approach, although the quickest means of *manual* recovery from a server failure, does not provide satisfactory availability of mission-critical data. Moreover, full-time on-site technical expertise is becoming a luxury few businesses can afford, particularly for distributed geographic locations. With the Compaq automated switchover process, the spare server becomes the active server and is back online in minutes—without administrator intervention.

This white paper provides the following information about the Compaq Standby Recovery Server:

- A general overview of the Standby Recovery Server, including software and hardware requirements, how the Standby Recovery Server works, and how long the Compaq automated switchover takes.
- Types of faults that can occur in a client server system and how the Standby Recovery Server detects and acts upon faults.
- Servicing a failed server and restoring the repaired server to operation.
- Effects of a server failure and switchover on clients.

The Appendix contains tables and a figure identifying system requirements.

Standby Recovery Server Automates Switchover

The Standby Recovery Server for Compaq ProLiant servers increases system availability for Microsoft Windows NT or NetWare environments by automating the process for detecting and replacing a failed server. The Standby Recovery Server allows one Compaq ProLiant server to act as a standby for another.

With the Standby Recovery Server, two ProLiant servers are both attached to a common set of ProLiant Storage Systems that holds a single copy of the operating system, applications, and data. (See Figure 1.) If the primary server fails, the ProLiant Storage System automatically switches from the primary to the recovery server. The recovery server then boots, and the system is back online in minutes.

Server Solution for Remote and Branch Offices

The Standby Recovery Server is designed for business-critical servers that cannot sustain long periods of downtime. It offers minimum downtime for customers with file or application servers at remote locations or branch offices where extensive technical expertise may not be available on site.

By automating the failure recovery process, the Standby Recovery Server eliminates the requirement for intervention by the system administrator to recover from a server failure. Thus, it increases system availability and offers an excellent solution for unattended, or “lights out,” operation. Moreover, the system administrator can schedule maintenance of the failed server at a convenient time, such as the end of the business day or a period of relatively light activity.

Benefits of Standby Recovery Server

The Standby Recovery Server offers customers the following benefits:

- Increased system availability as a result of fast server recovery
- Fully automated server switchover, ideal for “lights out” operation
- Ability to schedule service on a failed server at a convenient time
- Affordability, even for implementation at numerous operating locations
- Tower or rack-mounted server configurations

Key Terms

As used in this paper, the following terms have specific meanings:

Standby Recovery Server Two ProLiant servers attached to a common set of ProLiant Storage Systems with a single copy of the operating system, applications, and data. If the primary server fails, each ProLiant Storage System automatically switches from the primary to the recovery server. The recovery server then boots, and the system is back online in minutes without administrator intervention.

Recovery Server Option The option kit that includes the Recovery Server Switch and Recovery Server Interconnect.

Recovery Server Switch The intelligent SCSI switch installed in a ProLiant Storage System that automatically switches the SCSI bus between two servers.

Recovery Server Interconnect The RS-232 serial cable that connects the primary and recovery servers when the Recovery Server Option is used in the standby configuration.

primary server The primary server in the standby configuration that first boots the operating system upon power-up.

recovery server In the standby configuration, the recovery server listens to the primary server heartbeat.

Recovery Server Option Driver Software running in the primary server that generates the primary server heartbeat.

Compaq Recovery Agent Software contained in the system ROM BIOS that monitors the heartbeat transmitted by the Recovery Server Option Driver to the recovery server via the Recovery Server Interconnect. The Compaq Recovery Agent sends commands over the SCSI bus to the Recovery Server Switch(es) installed in the storage system(s).

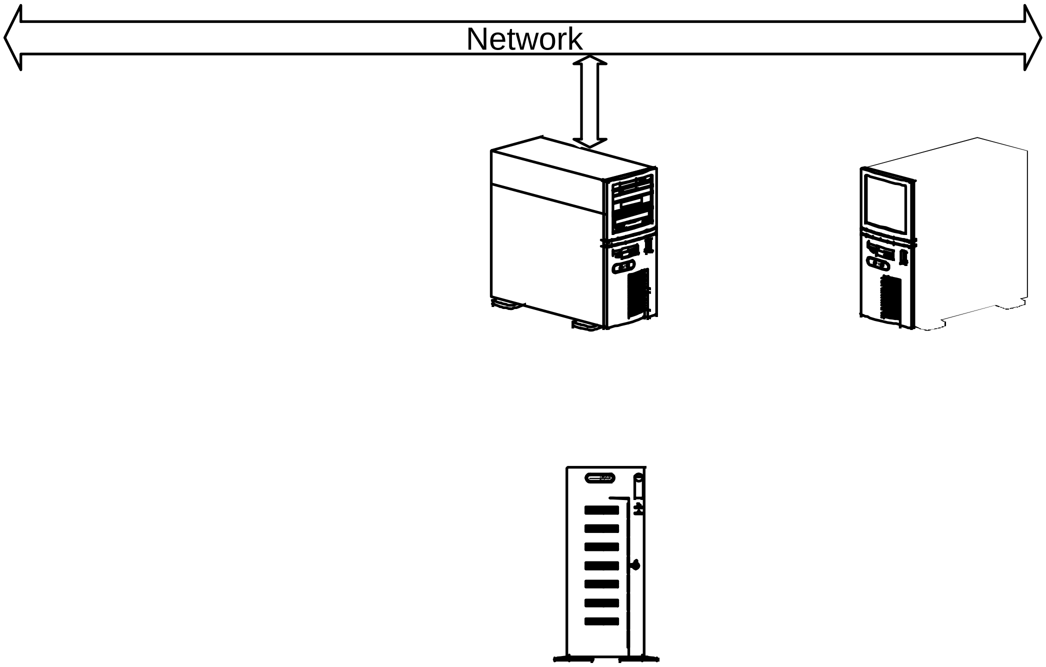
Network

SOFTWARE AND HARDWARE REQUIREMENTS

The Standby Recovery Server operates independently of the operating system and application software being used, but it requires the appropriate software drivers. Standby Recovery Server supports these operating systems only, on either tower or rack-mounted Compaq ProLiant Servers:

- Microsoft Windows NT 3.5X
- NetWare 3.12
- NetWare 4.X

The chance of data loss will be lowest with applications designed to behave predictably in the event of system failures. For example, most database applications use the concept of committed transactions to ensure data integrity if a system fails. Databases such as Oracle perform a database recovery operation when they are started. This recovery operation ensures the integrity of the database and removes transactions that had been neither committed nor completed when the system failed.

Applications that use a Compaq ProLiant server as a file server will operate the same way that they do on a single server. That is, with or without Standby Recovery Server, the same risk will exist for data loss caused by server failure.

For example, suppose that a server failure is caused by an operating system hang that fails to complete a disk write operation and then “locks up.” With Standby Recovery Server, the recovery server times out, switches the disks, and reboots. The data that were not written would be lost either with or without Standby Recovery Server.

The Standby Recovery Server supports these Compaq servers connected to Compaq ProLiant Storage Systems:

- ProLiant 4500
- ProLiant 4500R
- ProLiant 4000
- ProLiant 4000R
- ProLiant 2000
- ProLiant 2000R
- ProLiant 1500
- ProLiant 1500R

The Compaq ProLiant Storage Systems supported by the Standby Recovery Server are identified in Table 1 and the accompanying figure located in the Appendix of this paper.

With the Standby Recovery Server, the primary and recovery servers must have identical hardware configurations because they boot from the same Compaq ProLiant Storage System. In addition to identical memory configuration, the NICs and SCSI controllers must be installed in the same slots in both systems. Identical hardware configuration is required to prevent operating system and EISA (Extended Industry Standard Architecture) configuration problems when a switchover occurs and the recovery server boots from the external disk storage.

To allow for automated switchover from the primary server to the recovery server, all disk storage must be located externally in a ProLiant Storage System and must be attached to Compaq SMART array controllers. No internal disk storage can be used on either the primary server or the recovery server.

HOW STANDBY RECOVERY SERVER WORKS

The Standby Recovery Server increases system availability by automatically switching all disk storage from a failed primary server to the recovery server that is waiting to boot the operating system and reestablish access to data files.

The Standby Recovery Server requires two identically configured ProLiant servers connected by SCSI cables to a common set of ProLiant Storage Systems, which holds a single copy of the operating system, applications, and data. The primary and recovery servers are linked as shown in Figure 2 by the Recovery Server Interconnect, an RS-232 serial cable with specific pinout connections for use with the Standby Recovery Server.

The Recovery Server Option Driver runs on the primary server. There is a specific Recovery Server Option Driver for each operating system, so the correct driver must be loaded. Each server contains a SMART array controller. The electrical switching of the disk storage is performed automatically by the Recovery Server Switch, which is installed in each ProLiant Storage System connected to the two ProLiant servers.

Normal Operation

Figure 2 illustrates normal operation of Standby Recovery Server. Both the primary and recovery servers are attached to the same network. The primary server supports clients attached to it via the network, and the recovery server is idle.

Under normal operation, as soon as the recovery server has completed its power-on self test (POST) sequence, it executes the Compaq Recovery Agent contained in the system ROM BIOS. The function of the Recovery Agent is to monitor a periodic heartbeat message transmitted by the Recovery Server Option Driver to the recovery server via the Recovery Server Interconnect. At this point, no operating system is loaded on the recovery server. Thus, the recovery server is electronically attached to the network but is not accessible via the network. Its only function is to wait for the primary server to fail.

Receipt of the heartbeat message within a configured time-out period indicates that the primary server is functioning properly. The recovery server responds to each heartbeat with an acknowledgment message across the serial connection. As long as the recovery server continues to receive heartbeats according to schedule, it will remain in the idle mode.

The Recovery Server Option is an extension of the Automatic Server Recovery (ASR) functions currently supported in ProLiant servers. For more information about ASR, refer to the user’s guide provided with the server.

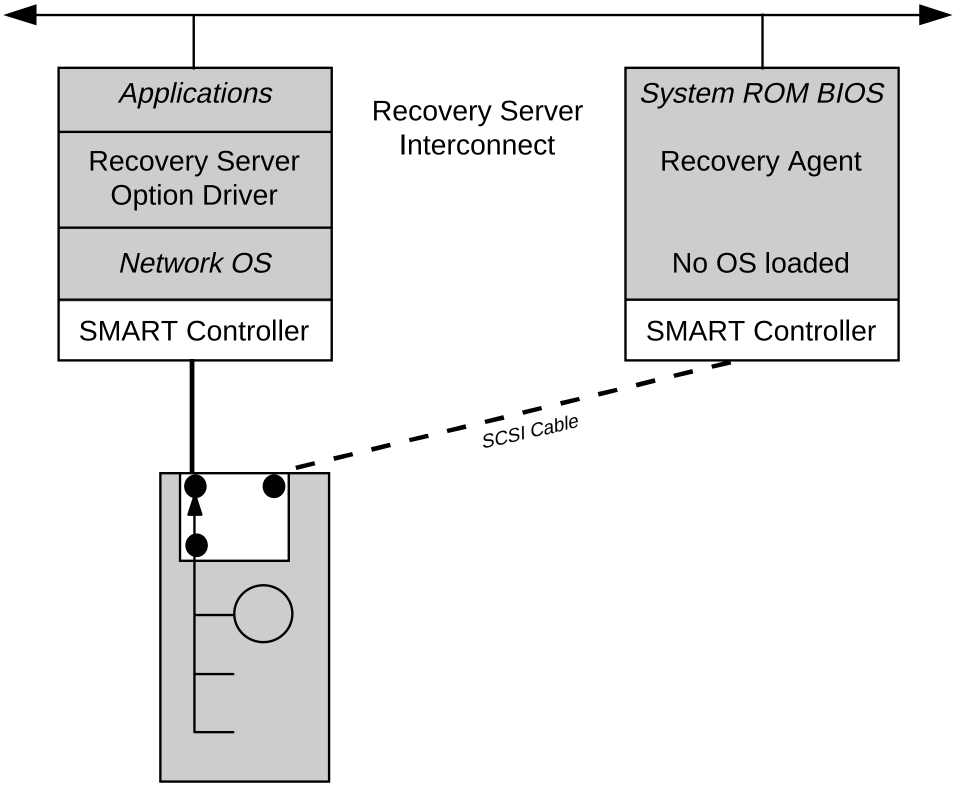
Applications
Recovery Server Option Driver
Network OS
SMART Controller
Recovery Server Interconnect
System ROM BIOS
Recovery Agent
No OS loaded
SMART Controller
SCSI Cable

Switchover Events

Switchover from the primary to the recovery server occurs when the primary server has failed. If the recovery server does not receive a heartbeat message within the time-out value set by the system configuration utility, the recovery server will assume that the primary server has failed. (Loss of the heartbeat message could occur either because the primary server has failed or because the connection of the Recovery Server Interconnect cable has been broken.)

The switchover events occur as follows:

1 . The Compaq Recovery Agent in the system ROM BIOS sends commands over the SCSI bus to the Recovery Server Switch installed in the common set of ProLiant Storage Systems. These commands cause the switch to disconnect the storage drives electrically from the primary server and then to connect them electrically to the recovery server.

2 . The recovery server proceeds through a normal boot sequence using the disk storage that was previously attached to the primary server.

3 . Because the servers are identically configured, when the boot process is completed, the recovery server assumes the logical network identity that was previously held by the primary server.

4 . At this point, network clients can log on to the recovery server and regain access to the data files. The specific actions required of clients will depend upon what they were doing at the time of server failure and upon the application that was in use.

With the Compaq automated switchover process, the recovery server becomes the active server and is back online in a matter of minutes—without administrator intervention.

Figure 3 illustrates a standby recovery server configuration after the switchover has occurred. The recovery server has assumed the function of the primary server. The primary server has completed an ASR reboot and is waiting to be serviced. The effects of server failure and switchover on clients are discussed later in this paper, in the section entitled “Client Behavior.”

For more information about the ASR reboot, refer to the user’s guide provided with the server.

Power Loss to Both Servers

If power is lost to both servers in a Standby Recovery Server configuration, the primary server will *not* boot in an unattended manner when the power is restored. An external power failure of this type will be recorded in the primary server NVRAM as a server failure requiring service, not as a power outage. Thus, when the primary server is powered on, the administrator will be prompted to run diagnostics or to press F8 to continue a normal boot sequence. This illustrates the importance of an uninterruptible power supply.

If the system is unattended when the power is restored, the recovery server will time out, switch the storage disks, and boot from the disks because the primary server is not sending the heartbeat message to the standby server.

SPEED OF THE SWITCHOVER

The Standby Recovery Server is designed for business-critical servers that cannot sustain periods of downtime exceeding several minutes. The total time required for the recovery server to assume the function of the primary server is predictable. It is the sum of the following six factors:

1 . The time that elapses from the moment at which a failure occurs in the primary processor to the moment at which that failure manifests itself in the loss of a heartbeat message. This time period may be very short (mere seconds) in the case of catastrophic failures such as loss of the processor, or it may be relatively long (several minutes) in the case of certain applications software failures.

2 . The defined time-out period that the Recovery Agent in the system ROM BIOS will wait for a heartbeat message before initiating a switchover is the ASR time-out value. It is set in the system configuration. The default value is 10 minutes. Available values range from 5 to 30 minutes.

3 . Once a switchover has been initiated, the time required to initialize the Compaq SMART array controllers and begin the operating system boot process from the drives, which by this time are electrically connected to the recovery server.

4 . The time required for the operating system to boot. This may be dependent upon the size and number of disk drives that are attached.

5 . Once the operating system is active, the time required for applications such as databases to be started and to perform integrity checks on their files. This is application and data dependent.

6 . Client log in.

To illustrate the effect of these factors, consider the example of a Standby Recovery Server configuration containing Compaq ProLiant 4000 servers with one SMART array controller and two 1GB drives operating with NetWare 3.12. The recorded times for an automated switchover event of this configuration are illustrated in Figure 4.

FAULTS

Server operation can be affected by multiple factors. In the Standby Recovery Server configuration, it is possible for several types of faults to occur, not all of which are primary server faults:

- Failures of the primary server, the types of faults for which the Recovery Server Option was designed.
- Loss of heartbeat resulting from serial cable problems, not from server problems.
- Failures that affect operation of the primary server but that do not cause a switchover.

Elapsed Time for an Automated Switchover

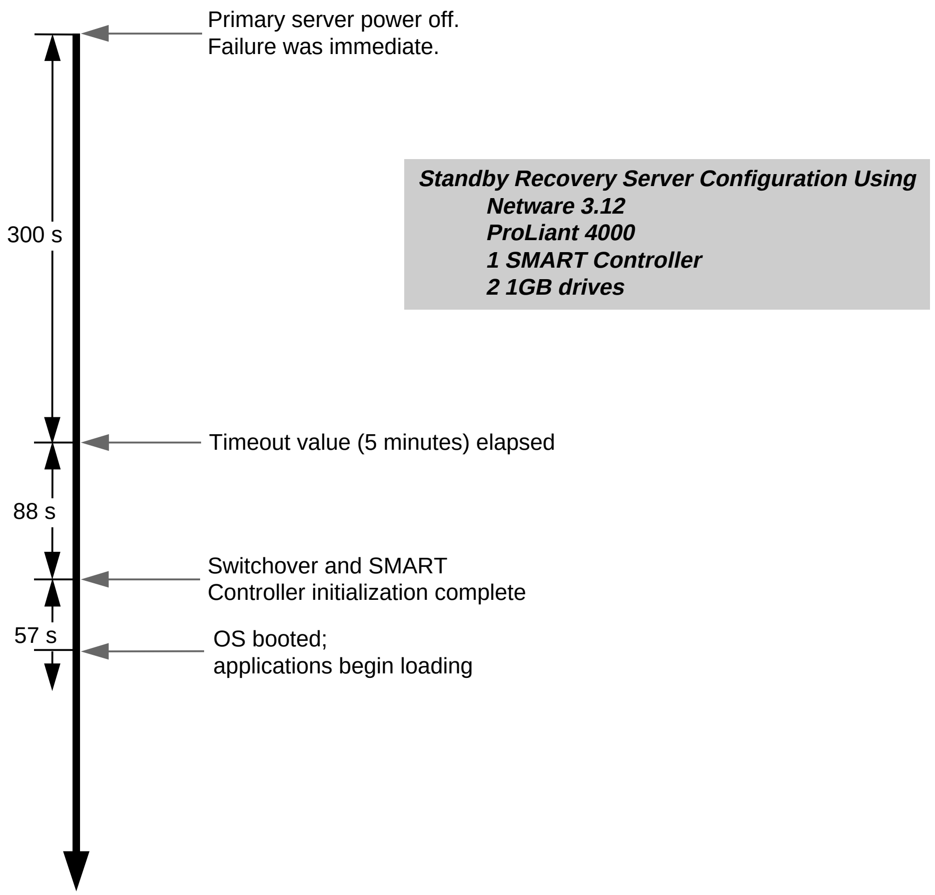
Primary server power off.
Failure was immediate.
300 s
Standby Recovery Server Configuration Using
Netware 3.12
ProLiant 4000
1 SMART Controller
2 1GB drives
Timeout value (5 minutes) elapsed
88 s
Switchover and SMART
Controller initialization complete
57 s
OS booted;
applications begin loading

Figure 4: In this example of an automated switchover event, the recorded elapsed time—from failure of the primary server through loading of the operating system on the standby server—totaled 445 seconds (less than 7½ minutes). Additional time is needed for loading the applications being used and for client log on.

FAILURE DETECTION

The failure detection mechanism in the Standby Recovery Server is based on the Recovery Server Option Driver software that runs in the primary server. As long as the recovery server receives the heartbeat message within the time-out period, it assumes that the primary server has not failed. Any failure in the primary server that stops the Recovery Server Option Driver from generating the periodic heartbeat message will be a detectable failure. Examples of detectable failures include:

- Catastrophic and unrecoverable hardware failure in the primary server such as loss of the processor or uncorrectable memory errors.
- Loss of the primary server power supply.

Generally, any failure that is detected by ASR will be detected and acted upon by the recovery server.

There is a class of failures that will cause the primary server to malfunction without causing loss of the heartbeat message. For example, failure of the NIC could render the primary server unusable, but the Recovery Server Option Driver would still send the heartbeat message to the recovery server. Failures of this type cannot be detected by the recovery server; therefore, an automatic switchover will not occur. Generally, the failures detected by the recovery server are the same ones that are detected by the ASR mechanism.

POTENTIAL INTERCONNECT FAILURES

The Recovery Server Interconnect may experience three types of failures. These failures and the behavior that they can cause are described below. This discussion assumes that Compaq Insight Manager is being used.

Primary Server Cable Failure

If the Recovery Server Interconnect is disconnected from the primary server, the recovery server cannot receive the heartbeat message. The Recovery Server Option Driver in the primary server can detect this condition. It sends an Insight Manager alarm indicating that the primary server has detected a cable fault and that it is shutting down the operating system in anticipation of the switchover that will occur because the recovery server is no longer receiving the heartbeat message.

Recovery Server Cable Failure

If the Recovery Server Interconnect is disconnected from the recovery server, the recovery server detects this condition and does not attribute loss of the heartbeat message to failure of the primary server. Because the primary server will no longer be receiving the acknowledgment message from the recovery server, however, the primary server will send an Insight Manager alarm indicating possible failure of the recovery server.

Severed Cable

If the Recovery Server Interconnect is physically cut, the heartbeat message and the acknowledgment message cannot travel between the primary server and the recovery server. Loss of the acknowledgment message will cause the primary server to send an Insight Manager alarm indicating possible failure of the recovery server. Meanwhile, loss of the heartbeat message for longer than the time-out period will cause the recovery server to switch the storage disks from the primary server to the recovery server and then boot.

Upon failure, the primary server normally becomes totally inactive. In the case of a severed cable, however, the primary server is not shut down. It continues running after its connection to the storage disks has been lost and the recovery server has booted. The operating system on the primary server can no longer function correctly, but the network protocol portion of the operating system is still active.

When the recovery server boots, it presents the same network identification as that being used by the primary server. As a result, network clients may not be able to log in to the server because both the primary and recovery servers are using the same network identification.

This type of failure is unlikely and is preventable with simple precautionary steps to protect the serial cable and its connections. Screw the serial cable down securely; and for maximum cable protection, rack mount the servers.

SERVICING THE FAILED SERVER

To reestablish Standby Recovery Server operation after a switchover, it is essential that a failed primary server be repaired or replaced and be brought back online. The Standby Recovery Server makes it possible for the system administrator to schedule service on the primary server at a convenient time while the recovery server is active. The primary server can be serviced on site or off site.

Once the switchover has occurred, no drives are electrically attached to the disk controllers in the primary server. For this reason, there may be some constraints on diagnostic activities that can be performed on the failed primary server on site. However, by disconnecting the primary server from the Recovery Server Interconnect and adding other drives to the primary server, full on-site diagnosis can be performed on the failed primary server while the recovery server is running, The failed primary server can also be disconnected from the recovery server and the ProLiant Storage System and be moved off site for service.

RESTORING THE CONFIGURATION AFTER SWITCHOVER

After the primary server has been either serviced or replaced, restore the original configuration. The recovery server and ProLiant Storage Systems must be power cycled in order to reinitialize the Recovery Server Switch. The disk drives will be electrically connected to the primary server and it will boot the operating system. The recovery server will return to its role of listening for the heartbeat message from the primary server.

After setting up and configuring both the primary and recovery servers, verify that both servers operate correctly and will switch over when needed.

CLIENT BEHAVIOR

When a failure of the primary server occurs, clients attached to the network experience a service outage. The length of this outage is described above. The symptoms experienced by the client vary depending on the operating system, the network protocol, and the application.

Windows NT Client Behavior

At a Microsoft Windows NT client, failure of the primary server appears as the inability to read from or write to the network device. In most cases, after the recovery server has booted, client connections to the recovery server will resume automatically.

Windows NT Event Log Entries

Disk Volume Configuration for NetWare 3.12

- In NetWare 3.12, it is not possible to configure the VREPAIR NLM to execute automatically before the SYS volume is mounted. Therefore, if the SYS volume is damaged by a primary server failure, NetWare 3.12 will not be able to mount the SYS volume. Thus, the SYS volume must be repaired by manually running the VREPAIR utility before NetWare can become operational and mount the SYS volume.
- Manual intervention will be required to run the VREPAIR utility to repair the SYS volume. However, commands can be included in AUTOEXEC.NCF to run VREPAIR automatically to process other NetWare volumes besides SYS, after SYS is mounted successfully.
- To reduce the possibility of damage to the SYS volume, Compaq recommends configuring the SYS volume primarily for use as storage for executables, not for data files. Executables are primarily read—not written to—so the chance of data loss is reduced.

Appendix—System Requirements

Table 1 lists all ProLiant Storage Systems and indicates that all except the original ProLiant Storage System are supported by the Recovery Server Option Kit. The illustrations in the figure below Table 1 are provided to assist in visual identification of tower ProLiant Storage Systems.

Table 2 summarizes system requirements for the Standby Recovery Server.

Table 1: Compaq ProLiant Storage Systems

North American Part No.	International Part No.	Chassis	Interface	Supported by Recovery Server Option
146700	146750	tower	Fast SCSI-2	no
197100	197150	tower	Fast SCSI-2	yes
189600	189640	tower	Fast-Wide SCSI-2	yes
163750	163755	rack-mountable	Fast SCSI-2	yes
189900	189905	rack-mountable	Fast-Wide SCSI-2	yes

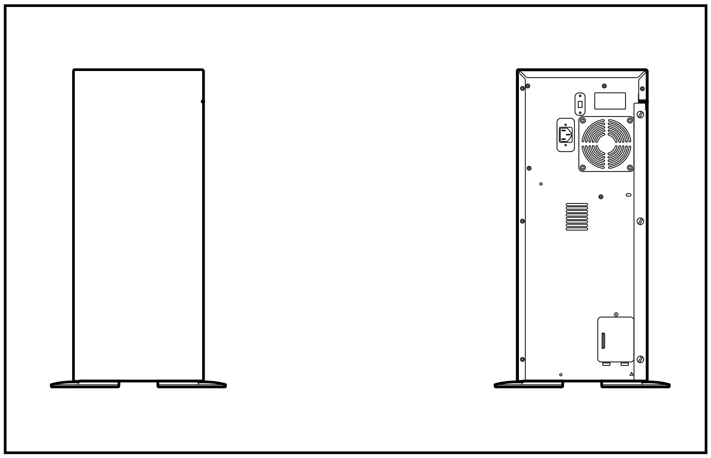

Table 2: System Requirements for the Standby Recovery Server

System Component	Requirement	Installation Notes
network operating system	One of these: • Microsoft Windows NT 3.5X • NetWare 3.12 • NetWare 4.X	
application software	Supports any.	Application software designed to behave predictably upon system failure will have less chance of data corruption as a result of a service outage. Oracle is a prime example.
Recovery Server Option Kit	One for each ProLiant Storage System used.	See *The Recovery Server Option User Guide* for the contents of this kit and installation instructions.
servers	Two each of one of these models: ProLiant 4500, 4500R, 4000, 4000R, 2000, 2000R, 1500, 1500R	The two servers must have identical hardware configurations.
SMART array controllers	Each SMART array controller can support up to two ProLiant Storage Systems.	The Array Accelerator on the SMART array controller must be disabled.
internal hard drives in server	None supported.	For internal CD-ROM and tape drives, integrated controllers may be used.
COM port	One serial port, COM Port 1 or 2, on each server for communication between primary and recovery servers.	The same COM port must be used on both the primary and the recovery servers.
external SCSI cables	One standard-to-wide cable required to connect each primary server and each recovery server to each storage system.	See Appendix B in *The Recovery Server Option User Guide* for cabling requirements.
external disk storage	Any ProLiant Storage System except Compaq part number 146700 (international part number 146750). Refer to Table 1.	All disk storage must be located in an external storage unit. Each external storage unit must have one Recovery Server Switch installed.

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, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, registered United States Patent and Trademark Office.

Aero, Concerto, QuickChoice, ProSignia, Systempro/XL, Net1, SilentCool, LTE Elite, Presario, SmartStation, MiniStation, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, TriFlex Architecture and UltraView, CompaqCare and the Innovate logo, are trademarks and/or service marks of Compaq Computer Corporation.

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

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

.

Compaq Standby Recovery Server

First Edition (May 1995)
Document Number 180A/0495