COMPAQ
Technical Guide

March 2001
14AL-0301A-WWEN

Prepared by Enterprise Appliance Business Unit

Compaq Computer Corporation

Contents

Disaster Recovery, A Planning and Recovery Guide for the TaskSmart N-Series Cluster

Abstract: This guide is provided to assist Compaq customers in employing disaster recovery plans and procedures for the Compaq TaskSmart N-Series Cluster. The objectives of this guide are to help customers:

- Understand the benefits of planning for a disaster and the importance of thorough documentation.
- Understand how to recover from a disaster based on several possible scenarios.

Some of the information in this technical guide is adapted from the *Compaq TaskSmart N-Series Cluster Installation Guide* and the *Compaq TaskSmart N-Series Cluster Planning Guide* that are available with the purchase of each Compaq TaskSmart N-Series Cluster appliance.

Introduction

Disaster recovery technology is an option that is essential for any business that is concerned about saving time and money. Whereas the primary purpose of backup technology is to protect and restore data and software applications, the primary purpose of disaster recovery technology is to restore the operating environment and data quickly. A good backup solution has disaster recovery technology intelligent enough to recover the server quickly and to exacting detail, returning it to its state prior to the disaster. This method is called "point-in-time" recovery. Simply stated, after the recovery process is complete, the server is returned to a state equal to what it was at the point in time of the last successful backup session.

A good, comprehensive backup solution is one of the most important aspects to a disaster recovery solution. The most obvious reason for employing a disaster recovery solution is that computers do break down. When this event occurs, mission-critical information becomes unavailable. A good backup solution will ensure that the information stored on a computer will be available when it is needed. When a network server fails due to human error, hardware failure, or a major disaster, the system must be carefully recovered before the applications and backed-up data can be restored. Disaster recovery technology strategically complements backup/restore technology. Backup/restore procedures require that a computing environment exists that will support data recovery procedures. Disaster recovery ensures that the environment is available and minimizes the amount of time required to bring network systems back to full functionality.

Microsoft Windows 2000 has several components that must be backed up together. These components make up the System State. The restoration of the System State, which must replace boot files first and commit the system hive of the registry as a final step in the process, is critical to the system recovery. In addition to restoring the System State, it is crucial to restore and recover the Quorum Disk. Recovery of the Quorum Disk re-establishes the clustering to its state prior to the disaster.

Note: This document uses the term **quiesce**. The term quiesce refers to a temporarily inactive state of a resource.

Determining Your Backup Solution

When designing the appropriate backup and restore solution, several key areas must be considered. Important issues to consider are backup and restore performance, media reliability, tape rotation schemes, and offsite storage of data. These topics are addressed further in this discussion, but this document will first focus on three key considerations. It is assumed at this point that the importance of protecting information is recognized, and a Compaq integrated solution has been chosen to fill this need. Other core issues must be considered when a backup solution is deployed. These issues include several key factors that must be addressed:

- Hardware
- Software (backup software vendor)

In many departmental and workgroup situations, stand-alone tape drives and tape libraries are connected directly to the server appliance, providing the server with fast backup speeds and the exclusive use of the devices. In enterprise situations, multiple servers commonly share a large tape library device through the existing communications network. This provides the benefit of an automated solution for reduced human error without impacting network performance during the

backup and restore process. This is the backup solution that will be supported for the Compaq *TaskSmart*™ N-Series Cluster.

Depending on the need or ability to share a tape library, a storage area network (SAN) can be deployed for backup and restore. This allows multiple servers to share an automated tape library over a dedicated communication network using Fiber Channel as the interconnect. This solution takes advantage of the Fibre Channel infrastructure, and provides greater flexibility in the distance between devices (up to 10 km using Single-Mode Fibre Channel), as well as greater speed (100 MB/s). The use of a SAN also provides better asset utilization, because multiple servers can share a larger automated tape library.

Hardware

A Modular Data Router (MDR) is required to connect the tape device to a fibre environment. This MDR provides a SCSI-to-Fibre bridge between the tape device and the fibre switch. Selecting the correct tape device and connection type ensures a reliable backup of data that is well suited to your particular computing environment. Figure 1 illustrates the cabling scheme for the Modular Data Router.

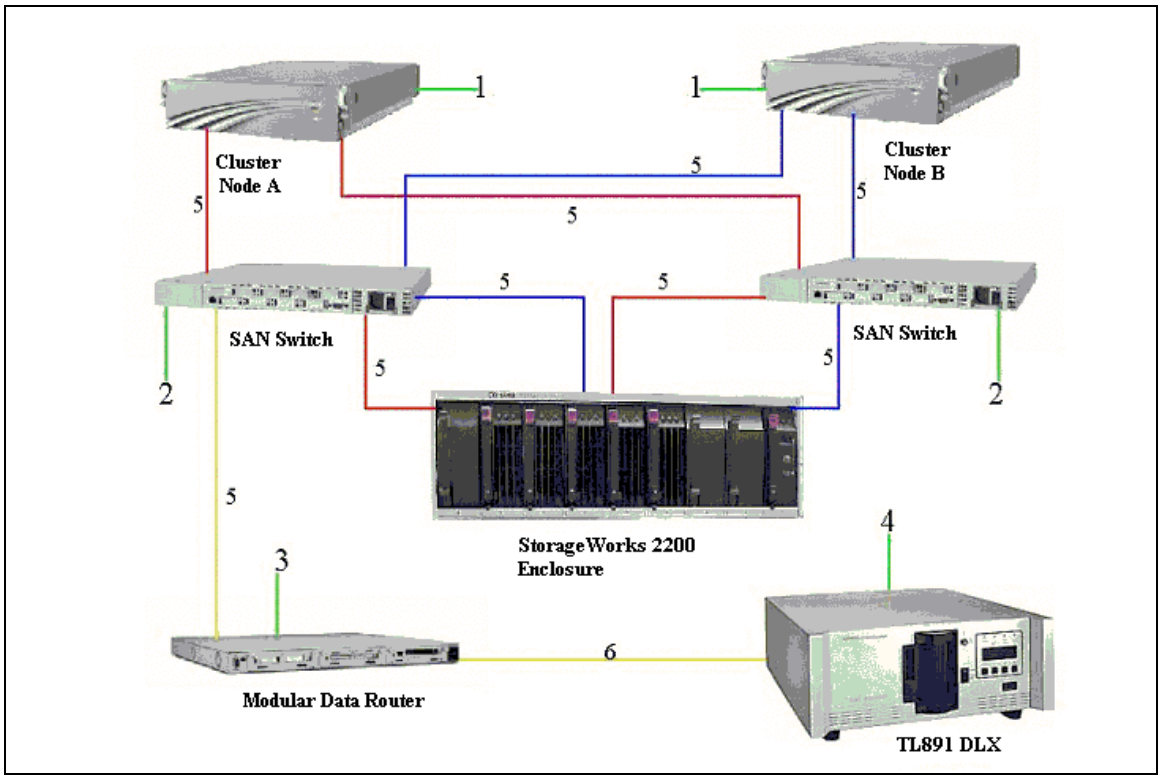

Figure 1. Modular data router cabling scheme

Table 1. Modular data router cabling scheme components

Item	Description
1	TaskSmart N-Series Clustering Node
2	Compaq Fibre Channel Storage Hub 12
3	Compaq Modular Data Router
4	Differential SCSI Tape Library
5	Multi-mode Fibre Channel Cable
6	External VHDCI SCSI Cable

Tape Device Type

Because their common usage in the industry, excellent performance, and capacity, Compaq recommends the following tape solutions for use with the TaskSmart N-Series appliance. All of these solutions are based on the Compaq DLT 35/70 (also known as DLT7000) tape drive.

- Compaq DLT Tape Array
- Compaq TL891 DLT MiniLibrary
- Compaq TL895 DLT Library
- Compaq ESL9326 DLT Library
- Compaq SSL2020TL AIT Library
- Compaq TL891DLX DLT Library
- Modular Data Router (SCSI-to-Fibre bridge)

These devices have been tested successfully with the TaskSmart N-Series appliance. Other tape libraries will be qualified. For a full list of qualified tape solutions, refer to:

http://www.compaq.com/TaskSmart/n2400

Switched Fabric Fibre Channel

Fibre Channel Switched Environments use a host bus adapter (HBA), and a fiber switch. A fibre switched environment is more scalable in performance and capacity than loop environments because the architecture provides a dedicated data path between two devices in a switch. Other devices may be connected to the switch device, but they do not see or interfere with the two nodes that are communicating.

This environment provides much greater performance, because loop-based environments must share a 100-MB/s data path. Switched environments contain many dedicated 100-MB/s data paths, without sharing the bandwidth.

The TaskSmart N-Series Cluster Fibre Channel Switched Environment will consist of:

- Two TaskSmart N2400 servers, each using two KGPSA-CB PCI to Fibre HBAs
- Two Fibre Channel 8- or 16-port *StorageWorks*™ SAN Switches

- Fibre Channel-to-SCSI bridge for library connectivity (Modular Data Router)
- Compaq TL891 MiniLibrary, TL895 DLT Library or ESL 9326DLT Library or any Compaq Tape Library that is supported on the TaskSmart N-Series Cluster

Software

Backup software vendors often offer a disaster recovery product addition that is used in conjunction with a backup solution. The supported software packages for the TaskSmart N-Series Cluster are:

- Veritas Backup Exec v 8.5
- Legato NetWorker v 6.0

These two software packages and their disaster recovery solutions are detailed in this document. Additional information for both of these software packages can be found at the following links:

Legato:

http://www.legato.com/products/protection/networker/networker6/

VERITAS:

http://www.veritas.com/us/products/backupexec/

Recovering from Controller Failure

Documentation is a crucial part of any disaster recovery solution. The information that is extracted from the HSG80 controller should be current and stored in a location separate from the servers it represents. In the event that both controllers fail during a disaster, thorough documentation will be a key to quick recovery. This section will demonstrate how to extract the HSG80 configuration into a text document and what information should be extracted.

In Hyper Terminal:

Select **Transfer**, and then select **Capture Text**.

Next, select the location where the file will be saved. It is recommended that this file is not saved on any of the disks managed by the controllers. Best practice would be to save it in a remote location.

The following are the recommended commands to type to extract the output to the specified file. When these commands are executed, the output will be saved in the text file specified.

Show This	This command will display the current configuration of the controller that has an established connection.
Show Other	This command will display the current configuration for the secondary controller.
Show Devices	This command will display all disks attached to the HSG80 controllers, their port, target and LUN, and what they are members of.
Show Units	This command will display the LUNs on the controllers and what disksets are associated with them.

Show Connections	This command will display the connections to the controllers and their unit offset.

Some optional commands can also be used to display further information:

Show Mirrorsets	This command will display all mirrorsets configured on the controllers along with the associated disks.
Show RAIDsets	This command will display all RAIDsets configured on the controllers along with the associated disks.
Show Sparesets	This command will display all sparesets configured on both controllers along with the associated disks.

Single-Controller Failure

In the event that a single controller fails, the recovery process is simple. Because the Compaq StorageWorks enclosure model 2200 is configured with redundant controllers, the failure of one controller will not result in data loss.

IMPORTANT: New controller hardware must be compatible with the remaining controller hardware before replacement procedures begin. See the product-specific release notes that accompanied the software release for information regarding hardware compatibility. The software versions and patch levels must be the same on both controllers. The new cache module must contain the same memory configuration as the module being replaced.

Replacing a Controller and Cache Module in a Dual-Redundant Controller Configuration

Use the steps in the following sections, "Removing a Controller and Cache Module in a Dual-Redundant Controller Configuration" and "Installing a Controller and its Cache Module in a Dual-Redundant Controller Configuration," to replace a controller and its cache module. Both cache modules must contain the same cache memory configuration.

Removing a Controller and Cache Module in a Dual-Redundant Controller Configuration

Use the following steps to remove a controller and its cache module.

1. Connect a computer or terminal to the maintenance port of the operational controller. The controller connected to the computer or terminal becomes "this controller;" the controller being removed becomes the "other controller."
2. Disable failover with the following command:

   SET NOFAILOVER

3. Remove the program card electrostatic discharge (ESD) cover and program card from the "other controller." Save them in a static-free place for the replacement controller.
4. Start the field replacement utility (FRUTIL) with the following command:

   RUN FRUTIL

   Follow the onscreen instructions to remove the elements.

CAUTION: The device ports must quiesce before the controller is removed. Quiescing is indicated by an "All device ports quiesced" message. Failure to allow the ports to be quiesced might result in data loss. Quiescing might take several minutes.

Note: A countdown timer allows a total of four minutes to remove both the controller and cache module. After four minutes, "this controller" exits FRUTIL and resumes operations. If this happens, return to step 4 and proceed.

When instructed to remove the controller and cache module, use the following steps:

1. Disconnect all fiber cables from the "other controller." For cables without extender clips, use thin needle-nose pliers to disconnect each cable.
2. Disengage both retaining levers on the controller and remove the "other controller," and then place the controller in an antistatic bag or on a grounded antistatic mat.
3. Disengage both retaining levers on the controller and remove the "other controller" cache module, and then place the cache module on a grounded antistatic mat or an antistatic bag.
4. If a replacement controller and cache module are not available, press the **N** key. FRUTIL will exit. Then, disconnect the computer or terminal from the controller maintenance port.
5. If a replacement controller and cache module are available, remove the DIMMs from the "other controller" cache module for installation in the replacement cache module.
6. Press the DIMM retaining clips down at both ends of the DIMM being removed.

Note: To make pressing down on the DIMM retaining clips easier, consider using the eraser end of a pencil or a small screwdriver.

7. Gently remove the DIMM from the DIMM slot, and then place it in an antistatic bag or on a grounded antistatic mat.
8. Repeat step 6 and step 7 to remove each DIMM.
9. Insert each DIMM straight into the appropriate slot of the replacement cache module, ensuring that the notches in the DIMM align with the tabs in the slot.
10. Press the DIMM gently into the slot until seated at both ends.
11. Engage two retaining clips for the DIMM.
12. Make sure that both ends of the DIMM are firmly seated in the slot and both retaining levers engage the DIMM.
13. Repeat step 9 through step 12 to replace each DIMM.
14. Press the **Y** key.
15. When complete, exit out of FRUTIL.

Installing a Controller and its Cache Module in a Dual-Redundant Controller Configuration

Use the following steps to install a controller and its cache module:

1. Insert each DIMM straight into the appropriate slot of the replacement cache module, ensuring that the notches in the DIMM align with the tabs in the slot.
2. Press the DIMM gently into the slot until seated at both ends.
3. Engage two retaining clips for the DIMM.
4. Make sure that both ends of the DIMM are firmly seated in the slot and both retaining levers engage the DIMM.
5. Repeat step 1 through step 4 for each DIMM.
6. Connect a computer or terminal to the maintenance port of the operational controller. The controller connected to the computer or terminal becomes "this controller"; the controller being installed becomes the "other controller."
7. Start FRUTIL with the following command:

   ```
   RUN FRUTIL
   ```

   Follow the onscreen instructions to install the elements.

Note: A countdown timer allows a total of four minutes to install both the cache module and controller. After four minutes, "this controller" will exit FRUTIL and resume operations. If the countdown timer expires, return to step 7 and proceed.

CAUTION: Carefully align the cache module and controller in the appropriate guide rails. Misalignment might damage the backplane.

Note: When the replacement cache module and controller are fully seated, the replacement controller restarts automatically—the reset LED turns on.

---

8. Press the **Enter** key to continue. The "other controller" restarts and FRUTIL exits.

Note: When a controller restarts, a visual indication is the temporary cycling of the port LEDs and a flashing reset button.

If the "other controller" did not restart, complete step 9 through step 11:

9. Press and hold the "other controller" reset button.
10. Reseat the "other controller" program card.
11. Release the reset button.

Note: In mirrored mode, FRUTIL initializes the mirrored portion of the new cache module, checks for old data on the cache module, and then restarts all device ports. After the device ports restart, FRUTIL tests the cache module and the ECB. After the test completes, the device ports requiesce and a mirror copy of the cache module data is created on the newly installed cache module.

12. Replace the program card ESD cover.

IMPORTANT: If the controller being installed was previously used in another subsystem, purging the controller of the old configuration is required (see CONFIG RESET in the controller CLI reference guide).

13. Enable failover and re-establish the dual-redundant controller configuration with the following command:

    SET FAILOVER COPY=THIS_CONTROLLER

    This command copies the subsystem configuration from "this controller" to the new controller.

14. If desired, verify the failover configuration with the following command:

    SHOW THIS_CONTROLLER FULL

15. Configure the controller using the TaskSmart N-Series Cluster installation guide as a reference.
16. Connect all host bus cables to the new controller.
17. Disconnect the computer or terminal from the controller maintenance port.

Replacing a Controller in a Dual-Redundant Controller Configuration

Use the following steps in "Removing a Controller in a Dual-Redundant Controller Configuration" and "Installing a Controller in a Dual-Redundant Controller Configuration" to replace a controller.

Removing a Controller in a Dual-Redundant Controller Configuration

Use the following steps to remove a controller:

1. Connect a computer or terminal to the maintenance port of the operational controller.
2. The controller connected to the computer or terminal becomes "this controller"; the controller being removed becomes the "other controller."
3. Disable failover and take the controllers out of the dual-redundant configuration with the following command:

    SET NOFAILOVER

4. Remove the program card ESD cover and program card from the "other controller."
5. Save them in a static-free place for the replacement controller.
6. Start FRUTIL with the following command:

    RUN FRUTIL

    Follow the onscreen instructions to remove the controller.

CAUTION: The device ports must quiesce before the controller is removed. Quiescing is indicated by an "All device ports quiesced" message. Failure to allow the ports to quiesce might result in data loss. Quiescing might take several minutes. ESD can easily damage a controller. Wear a snug-fitting, grounded ESD wrist strap.

Note: A countdown timer allows a total of four minutes to remove the controller. After four minutes, "this controller" will exit FRUTIL and resume operations. If the countdown timer expires, return to step 6 and proceed.

When instructed to remove the "other controller," use the following steps:

7. Disconnect all fiber cables from the "other controller." For cables without extender clips, use thin needle-nose pliers to disconnect each cable.
8. Disengage both retaining levers and remove the "other controller," and then place the controller in an antistatic bag or on a grounded antistatic mat.
9. If a replacement controller is not available, press the **N** key and FRUTIL will exit. Then, disconnect the computer or terminal from the controller maintenance port.
10. If a replacement controller is available, press the **Y** key and then follow the onscreen instructions. Then, exit out of FRUTIL. This procedure will restart the other controller.

Note: When fully seated, the replacement controller restarts automatically—the reset LED turns on.

Installing a Controller in a Dual-Redundant Controller Configuration

Use the following steps to install a controller:

1. Connect a computer or terminal to the maintenance port of the operational controller.
2. The controller connected to the computer or terminal becomes "this controller;" the controller being installed becomes the "other controller."
3. Start FRUTIL with the following command:

   ```
   RUN FRUTIL
   ```

   Follow the onscreen instructions to install the controller.

Note: A countdown timer allows a total of four minutes to install the controller. After four minutes, "this controller" will exit FRUTIL and resume operations. If the countdown timer expires, return to step 2 and proceed.

4. Carefully align the controller in the appropriate guide rails. Misalignment might damage the back plane.

Note: When fully seated, the replacement controller restarts automatically—the reset LED turns on.

5. Press the **Enter** key to continue.

   The "other controller" restarts and FRUTIL exits.

Note: Restart indication is the temporary cycling of the port LEDs and a flashing reset button.

If the "other controller" did not restart, complete step 5 through step 7:

6. Press and hold the "other controller" reset button.
7. Reseat the "other controller" program card.
8. Release the reset button.
9. Replace the program card ESD cover.

IMPORTANT: If the controller being installed was previously used in another subsystem, purging the controller of the old configuration is required (see CONFIG RESET in the controller CLI reference guide).

---

10. Enable failover and re-establish the dual-redundant controller configuration with the following command:

    SET FAILOVER COPY=THIS_CONTROLLER

    This command copies the subsystem configuration from "this controller" to the new controller.

11. If desired, verify the failover configuration with the following command:

    SHOW THIS_CONTROLLER FULL

    See the controller TaskSmart N-Series Cluster Installation Guide to configure the controller.

12. Connect all host bus cables to the new controller.
13. Disconnect the computer or terminal from the controller maintenance port.

Dual-Controller Failure

If both nodes are unaffected by the disaster, shut them down and use a Hyper Terminal connection from another computer. It is best if the nodes are not up when configuration of the controllers takes place. It is extremely important that the controllers are configured exactly the way they were prior to disaster, the section titled "Recovering From Controller Failure" earlier in this guide details the process for documenting the controller settings. It will also be necessary to replace all damaged components. When both controllers have been replaced it is necessary to complete the configuration process that is explained in the TaskSmart N-Series Cluster installation guide. It is critical that up-to-date and complete information be present to provide quick and efficient recovery time. This section assumes the following conditions:

- Both nodes are in working condition and the cluster state is unharmed.
- The storage cabinets and the data residing on them are still intact.
- Current documentation of the controller configuration is on hand.
- A current backup of all data is on hand.

Use the following steps for recovery from dual controller failure:

1. Shut down both nodes.
2. Replace both controllers and their cache modules. (This procedure is documented in the section titled "Single Controller Failure" earlier in this paper)

CAUTION: Do **not** initialize the disksets as they are created. This will result in total data loss.

---

3. Reconfigure the controllers to their original configuration according to the documentation that was extracted from the controllers before failure. It is very important that all units, disksets, and any specific configurations pertaining to them are restored exactly to their previous state.

   Disable access to all units, except to the Quorum Disk unit using this command:

   ```
   Set d1 disable_access=all
   ```

   This command must be entered for all units, except for the Quorum Disk unit. For more information on enabling and disabling access to units, see the installation guide.

4. Restart Node A. Do not restart Node B at this time.

   Confirm that Cluster Services has started. This can be accomplished by opening Cluster Administrator and confirming that the cluster resources are online. The pool and share resources will be offline at this point. This is expected behavior.

5. Enable access to each unit, one at a time. After access has been enabled to the unit, rescan for new devices in device manager. Follow this procedure until access has been enabled for all units. To enable access to a unit, the following command is used:

   ```
   Set d1 enable=!newcon00,!newcon03,!newcon04,!newcon05
   ```

   For more information on enabling and disabling access to units, see the TaskSmart N-Series Cluster installation guide.

6. Restart Node A.
7. Confirm that all SANworks Virtual Replicator resources are online. Confirm this in Cluster Administrator and Snapshot Manager.
8. Restart Node B.
9. Confirm that Node B successfully joined the cluster.
10. In Cluster Administrator, do the following:
    - Move the cluster group to Node B.
    - Move a pool group to Node B.

Restoration of a Quorum Disk with NTBackup and the ClusRest Tool

This section covers the restoration of a Quorum Disk as well as the physical replacement of the hard disk. It is optimal for a Quorum Disk to be configured in a mirrorset. It is possible for a spareset to be created and used with the Quorum Disk.

Quorum Disk Replacement

The Quorum Disk can be backed up in the same manner as other data. However, it must be restored using a specific tool provided by Microsoft. This tool is used in the cluster after a node has been restored using NTBackup. NTBackup leaves the Quorum data in a directory on the disk of the node, but does not restore it to the Quorum Disk at that time. Restoring the Quorum Disk requires that the cluster be stopped and restarted, which may not be necessary or desirable. If a Quorum Disk requires a restore, the cluster is most likely down, and the operation will not negatively impact operation.

The process of restoring a single node of a cluster is straightforward. The System State must be restored, and then the cluster state is restored to the node. Then, the system must be restarted. This causes the node to rejoin the cluster and restores it to operation.

After restoring the node, run clusrest.exe to restore the Quorum Disk. This moves the Quorum data from the node to the Quorum. The node can then be restarted and the cluster is established. The other node may join the disk as it is brought up.

For information on where to obtain and install the ClusRest tool, refer to Appendix A.

CAUTION: Restoring the Quorum Disk rolls the cluster back in time to the backup date. There are impacts to performing this operation that can include loss of data. This operation should be undertaken only when it becomes absolutely necessary.

Physical Replacement of a Disk in a RAIDset or Mirrorset

This section assumes that the administrator has backed up the data before replacing disks. Disk replacement has several variations. Disk replacement depends on if the disk is online or failed. It also depends on the way the disks have been placed into disksets.

Scenario 1: Mirrorset or RAIDset with a Failed Disk

When a disk fails, the HSG80 controller marks that disk as failed and places it into a failedset. To confirm this, type Show Devices in the HyperTerminal console. If disk 10100 is failed, then it must be in a failedset.

Note: For this example: R1 is a RAIDset and disk10100 is a failed disk from R1.

1. Remove the failed disk.
2. At the Hyper Terminal console, type:
   - Show R1 (This command will show the mirrorset policy settings. Record this information.)
   - Delete Failedset Disk10100 (This command will delete the failed disk from the failedset)
3. Insert the new disk.
4. At the Hyper Terminal console, type:
   - Set R1 nopolicy (Disables any controller replacement mechanisms)
   - Run Config
   - Set R1 replace=disk10100 (This command will add diskxxxxx to the mirrorset)
   - Set R1 policy= (Enter the policy settings from first step)

Scenario 2: Mirrorset or RAIDset with a Failed Disk and a Spareset

Note: For this example, R1 is a RAIDset, disk10100 is a failed disk from R1, and disk11300 is a spareset.

1. Remove the failed disk.
2. At the Hyper Terminal console, type:
   - Show R1 (Will show the mirrorset policy settings. Record this information)
   - Delete Failedset disk10100
3. Insert the new disk.
4. At the Hyper Terminal console, type:
   - Run Config
   - Set R1 nopolicy
   - Set R1 remove=disk11300 (This will remove the spare disk that was added automatically into the set.)
   - Set R1 replace=disk10100 (This will add the original disk into the set.)
   - Delete failedset disk11300 (When the remove command is executed, the controller removes the specified disk and makes it a failedset, deleting the failedset.)
   - Add spareset disk11300 (This adds the original spareset disk into a spareset.)
   - Set R1 policy=best_performance (This sets the policy of the RAIDset back to best_performance or any other setting.)

Recovery from a Storage Enclosure Failure and Data Loss

This section covers the recovery procedures for the failure of a storage enclosure that results in the loss of more than one disk per RAIDset or mirrorset. If only one disk per RAIDset or mirrorset is lost, then refer to the section titled “Physical Replacement of a Disk in a RAID or Mirrorset” earlier in this paper.

This recovery assumes the following statements to be true:

- A current backup of the lost data is on hand.
- Documentation of the controller configuration is on hand.

If the disk devices in the storage cabinet are recoverable, then perform the following steps:

1. Shut down Node A and Node B.
2. Remove all disks from the destroyed or damaged storage enclosure.

3. Replace the storage enclosure and place the disk devices back into the storage enclosure. It is extremely important to place the disk devices back into the enclosures in the same position they were in previous to the disaster. If more than one enclosure has been lost, then it is necessary to repeat this procedure for all damaged enclosures one at a time.
4. Bring all disk enclosures back online.

   The controllers will query the disk devices and re-establish data resources.
5. Run a SHOW DEVICES command on the controller. Confirm that there are no failedsets from the mirrorsets or RAIDsets. (If a controller queries a disk that is failed, it will remove the disk from the diskset and place it into a failedset.) If failedsets are present, follow the instructions in the section titled "Physical Replacement of a Disk in a RAIDset or mirrorset" earlier in this paper.
6. Reboot both Node A and Node B.

If the disk devices in the storage enclosure are not recoverable, then perform the following steps:

1. Shut down Node B.
2. Document your current configuration in Snapshot manager. On Node A, delete all SANworks Virtual Replicator information that has been affected by the failures, including all pools, virtual disks, and snapshots.
3. Document your current configuration in Cluster Administrator. In Cluster Administrator, delete all resources that pertain to the pools, virtual disks, and snapshots that were deleted in the previous step, including the network name, IP address and file share resources that pertain to the deleted pool(s).
4. Shut down Node A.
5. Replace the storage enclosures, and then place the disk devices back into the storage enclosure. If more than one enclosure has been lost, then it is necessary to repeat this procedure for all damaged enclosures, one at a time.
6. On the controller, run the following commands:
   - Set no_failover
   - Run Frutil, press the **N** key, then select **4(I/O Module)**

   Follow the prompts until the process is complete.
7. Type Set multibus_failover copy=other
8. Delete the units, RAIDsets or mirrorsets of which the damaged disk devices were members. A unit must be deleted first, and then the diskset can be deleted.
9. Recreate all RAIDsets and mirrorsets. Initialize the disksets with the following command before adding them to a unit:

   INIT R1 (Where R1 is the name of the diskset)
10. Recreate all units and add the disksets as members of the appropriate units.
11. Disable the **All** option for port security:

    Set D100 disable=all (Where D100 is the unit number. Perform this command for all new units.)
12. Restart Node A.

13. Enable port security on each unit, one at a time: (This will allow sufficient time for the OS to configure the new devices.)
    a) Enable access from all connections for the Quorum Disk unit.
    b) Run command Set D100 enable= (Connection names, all units that will be using a unit offset of 0 should grant access to connections on port one. All units that will be using a unit offset of 100 should grant access to connections on port two.)
14. Recreate all *SANworks*™ Virtual Replicator resources that were deleted in the prior steps.
15. In Cluster Administrator, complete the following tasks:
    – When a pool is created in SWVR, an SE pool resource will automatically be created in Cluster Administrator. It will only be necessary to create the file share, network name and IP address resources for each group of which the pool(s) are members.
    – Recreate all resources that were deleted in prior steps. This may include the recreation of a virtual server and the group in which the pool resources reside.
16. Restore data from backup.

Recovery from a Single-Node Failure

If a single node fails within the cluster, there are several tools and procedures that can be used for a quick and clean recovery. The following section describes the steps necessary to perform a complete recovery and rejoin the cluster.

1. From Cluster Administrator, right click the node that has failed, and then select **Evict Node**.
2. Perform a QuickRestore on the failed node using the **Quick** option.
3. Configure the network settings.
4. Install cluster services:
    a) Select **Start**, **Settings**, and then **Control Panel**.
    b) Select **Add/Remove Programs**, and then **Add/Remove Windows Components**.
    c) Select **Configure** for Cluster Service.
    d) Select **Join As Second Node In A Cluster** to install cluster services.
5. Install SANworks Virtual Replicator:
    a) Insert the SANworks Virtual Replicator CD into the CD ROM drive.
    b) Choose **Complete Setup**.
    c) Upon completion, restart the node.
6. When the node comes online, open Cluster Administrator and confirm that the node successfully joined the cluster. Move a group or the cluster group over to the other node to confirm functionality.

7. Select **Start**, **Programs**, and then **Compaq SANworks Virtual Replicator** (SWVR) to launch Snapshot manager.

   The SWVR resources will be displayed. If they are not displayed, do the following:

   a) From the **Action** menu, select **Connect To Another Computer**.

   b) Specify the name of the cluster. This refreshes the SANworks Virtual Replicator information back in Snapshot Manager.

Recovery from Dual Node Failure With NTBackup

If both nodes of the cluster are lost, the following procedures can recover the nodes along with the Quorum Disk. These procedures will restore the cluster to its working state before the disaster. It is assumed that a current backup of the System State and the Quorum Disk is on hand to perform this recovery, and that no configuration changes have occurred since the backup. It is critical that the data backup and System State and Quorum Disk backup have the same time stamp.

CAUTION: If changes have been made to SWVR resources that are not in the current System State and Quorum Disk backup, then data loss can occur. Ensure that a current backup of the System State and Quorum Disk occurs frequently.

It is imperative that a current backup of the System State and the Quorum Disk are kept in an offsite location to be used in the event of a disaster.

1. Use the TaskSmart N-Series Cluster QuickRestore CD to restore both nodes. Select the quick option from the menu.
2. When the QuickRestore process is complete, reconfigure the cluster interconnect private network settings and the public network settings on Node A. It is not necessary for the cluster to rejoin the domain on Node A, because the System State restore will reconfigure this. If NIC teaming is installed and configured, then a restart will be necessary.
3. Install the ClusRest tool: (This tool is required to restore the Quorum Disk.)

   Note: The ClusRest tool is for use only with NTBackup.

   For more information on the ClusRest tool, refer to Appendix A.
4. Open NTBackup. A prompt is displayed, **Import Media Present**.

   Select **Allocate All Compatible Import Media To Backup**.
5. On Node A, restore both the System State and the Quorum Disk. A system restart is required upon completion.

6. When the restart is complete, log onto the domain, and then run the ClusRest tool:
   - Open a command prompt.
   - Change to the directory where the ClusRest tool was installed. The default directory is c:\program files\resource kit.
   - Run the ClusRest tool by entering, CLUSREST.
   - Press the **Y** key to continue.
   - Upon completion, close the command prompt, and then restart Node A.
7. Reinstall Compaq SANworks Virtual Replicator:
   - Insert the SANworks Virtual Replicator CD into the CD ROM drive.
   - Select **Repair** from the SWVR installation menu.
8. Open SANworks Virtual Replicator Snapshot Manager to ensure that all pools, vdisks, and snapshots are present.
9. Open Cluster Administrator.
10. Evict Node B by right clicking **Node B** and selecting **Evict**.
11. Power up Node B. Install and configure all network configurations, including domain configuration.
12. Install cluster services.
    a) Select **Start**, **Settings**, and then **Control Panel**.
    b) Select **Add/Remove Programs**, and then **Add/Remove Windows Components**.
    c) Click **Configure** for Cluster Service.
    d) Select **Join As Second Node In A Cluster** to install cluster services.
13. Install SANworks Virtual Replicator:
    a) Insert the SANworks Virtual Replicator CD into the CD ROM drive.
    b) Select **Complete Setup**.
14. Upon completion, restart the node. After restarting, open SANworks Virtual Replicator. If the resources are not displayed, complete the following steps:
    a) From the **Action** menu, select **Connect To Another Computer**.
    b) Specify the name of the cluster. This refreshes all of the SANworks Virtual Replicator information in Snapshot Manager.
15. Configure any appropriate changes in Cluster Administrator. Configurations that may need to be altered are: (Because Node B was evicted, and then rejoined the cluster, any configurations specific to Node B must be re-established.)
    - Preferred nodes for group resources.

      Note: Ensure that the pool resource within each group has both nodes as a possible owner

    - Failover or failback settings for each resource.

Recovery Using Legato Networker

This section illustrates how to recover a Microsoft Cluster Server (MSCS). Familiarity with MSCS concepts and operations is required. The following topics are documented in this section:

- Recovering a Failed Quorum Disk
- Recovering One Cluster Node
- Recovering Multiple Cluster Nodes

Note: For more information on other backup software, refer to the documentation for that software.

Prerequisites

To recover an MSCS host in a TaskSmart N-Series Cluster environment, the administrator must ensure that each of the following prerequisites are satisfied:

- The NetWorker client is installed on each cluster node.
- Backups that include the SYSTEM save sets (SYSTEM FILES, SYSTEM DB, and SYSTEM STATE) have been performed on a regular basis by a NetWorker server host in the same domain as the cluster nodes. This will help ensure that data is available from the NetWorker server for recovery to the desired point in time. (The cluster database is a component of the Windows System State, and as such is automatically included when the System State save set is specified for backup or recovery).
- During the recovery, the domain controller for the cluster is available to authenticate the node joining the cluster.

Recovering a Failed Quorum Disk Using Legato NetWorker 6.0

This section describes the procedure for recovering a failed Quorum Disk. The procedure assumes the following:

- This procedure assumes that both disks in the Quorum mirrorset are unusable.
- The Quorum Disk is designated for exclusive use by MSCS.
- The Quorum Disk resides in a mirrorset.
- A spare disk (identical in type to the Quorum Disk) is available to replace the failed disk.
- The new Quorum Disk is to be assigned the same drive letter as the failed disk. For related information, refer to Microsoft Knowledge Base article *Q172944: How to Change Quorum Disk Designation ID*.
- The failed disk devices have been replaced, and the unit and mirrorset for the Quorum Disk have been recreated.

How to Recover a Failed Quorum Disk

To recover a failed Quorum Disk, perform the following steps:

1. Close all instances of Microsoft Cluster Administrator.
2. Stop the cluster service on both nodes.
3. Use the Computer Management Services facility to access the **Cluster Service Properties** screen. Change the Cluster Service Startup Type to **Manual** on both nodes.
4. Shut down both nodes.
5. Replace the failed Quorum Disk. Refer to the "Physical Replacement of a Disk in a RAIDset or mirrorset" earlier in this document for more information.
6. Power up both nodes.
7. Format the new disk with the same file system, partitioning scheme, drive letter, and label as the failed Quorum Disk. Verify that it is identical on both nodes.
8. On one node, use the Computer Management Services facility to access **the Cluster Service Properties** screen.
9. Type -fixQuorum as a start parameter, and then start the service.
10. Use Cluster Administrator to rename the failed Quorum Disk to **RemoveMe**.
11. Create a new disk resource **Disk*x***: (where *x* is the drive letter of the old Quorum Disk). Place the new disk resource in the cluster group.
12. Bring this disk resource online.
13. Right click the cluster name, and then select **Properties**.
14. Select the **Quorum** tab and make the new drive the Quorum Disk.
15. Use the Computer Management Services facility to access the Cluster Service Properties page:
    a) Stop the service.
    b) Remove -fixQuorum as a start parameter.
    c) Start the service.
16. Start the cluster service on the other node.
17. Delete the **RemoveMe** resource.
18. Bring the cluster group online.
19. Use the Computer Management services facility to access the **Cluster Service Properties** screen. Change the Cluster Service Startup Type to **Automatic** on both nodes.
20. Reboot both nodes.

Recovering One Cluster Node Using Legato NetWorker 6.0

This section describes how to restore cluster services in the event one of the cluster nodes fails. Corruption of a cluster-critical file may cause a partitioned cluster. This results in one node of the cluster being unaware of the presence of the other operating node. In this situation, each node may attempt to take control of the shared Quorum device, potentially rendering one node unable to function as a member of the cluster. In the following scenario, Node B has failed and the Quorum Disk has successfully failed over to Node A.

How to Recover One Cluster Node

To recover the cluster configuration, perform the following steps:

1. Using Cluster Administrator on Node A, evict Node B from the cluster.
2. Using the TaskSmart N-Series Cluster QuickRestore CD, restore Node B using the **Quick** option
3. Reinstall and configure NIC teaming if applicable.
4. From Node B, log on as Administrator to the domain in which the cluster nodes reside.
5. On Node B, select **Control Panel**, **Add/Remove Programs**, **Add/Remove Windows Components**, and then select **Configure** for cluster service.
6. During the wizard setup, select **Join An Existing Cluster**, and then enter the cluster name. After the wizard finishes on Node B, go to Cluster Administrator on Node A. From Cluster Administrator on Node A, ensure that Node B is available for failover operations.
7. Reinstall SANworks Virtual Replicator on Node B.
8. Only Node A, the node that owns the shared resources, may be running at the time of recovery of the cluster database. Therefore, stop the cluster service on Node B or shut it down.
9. From the **NetWorker User Recovery** window on Node A, mark at least the System State save set for recovery. The cluster database is a component of the Windows 2000 System State, and as such is automatically included when the System State save set is specified for backup or restore.

   Note: With NetWorker 5.7 (and later) for Windows 2000, the three System save sets are interdependent. It is recommended to restore all of the System save sets whenever you need to restore any particular System save set. For more information, refer to the section on save set interdependencies in the NetWorker administrator's guide, Windows Version.

10. Click **Start** to begin the System State recovery process.
11. After the System State save set is restored, reboot Node A.
12. Using Cluster Administrator on Node A, confirm that the cluster resources were restored to the point in time when the backup occurred. If a regularly scheduled backup has been performed using NetWorker, this will recover the cluster database to a point in time shortly before the loss of Node B.

13. Start the cluster service on Node B. To do so, at Node B, run the command `net start clussvc` from a command prompt, or select **Computer Management**, and then **Services**.
14. From Cluster Administrator on Node A, monitor the cluster joining status of Node B.
15. From Node B, use Cluster Administrator to verify that the cluster group can be moved between the nodes by right clicking the group, and then selecting **Move Group**. The cluster configuration should now have been recovered.

Recovering Both Cluster Nodes Using Legato NetWorker 6.0

This section provides general guidelines for performing a cluster recovery in the case of a failure of both cluster nodes. In this scenario, the operating system is unusable on each node. Therefore, this recovery procedure includes a QuickRestore of each node, as well as NetWorker recovery of the cluster database from a previous backup.

IMPORTANT: Because cluster configurations may vary, it is not possible to provide cluster disaster recovery procedures for every situation. Depending on the particular cluster configuration and the nature of the failure, it might be necessary to vary some of the procedures described in this section.

How to Recover Multiple Cluster Nodes

To perform a complete cluster recovery in a situation in which both cluster nodes, Node A and Node B, have failed:

1. Perform a QuickRestore on each node, by selecting **Full**. During the QuickRestore setup, verify that the domain controller is available and that each potential node is able to join.
2. Delete the MSCS folder on the Quorum drive.
3. On the Quorum Disk, run `chkdsk`.
4. Shut down Node B, start the MSCS installation on Node A from **Control Panel**, **Add/Remove Programs**, and then **Add/Remove Windows Components**. Select **Configure** for Cluster Services. The MSCS Cluster Wizard is displayed and provides guidance through the setup process.
5. During the setup process, using Cluster Wizard, enter the same configuration information that was used prior to the failure of the cluster nodes (including user account, IP addresses, and cluster name).
6. Reboot Node A.
7. Install MSCS on Node B, joining A.
8. Reboot Node B.
9. On Node B, run the `net stop clussvc` command from a command prompt.
10. Install the NetWorker client software on Node A.
11. From NetWorker User on Node A, select the System Files and System State save sets to recover to the desired point in time prior to the cluster failure. Click **Start** to begin the recovery process.
12. After the recovery on Node A is complete, reboot Node A.

13. Reinstall Compaq SANworks Virtual Replicator:
    a) Insert the SANworks Virtual Replicator CD into the CD ROM drive.
    b) Select **Repair** from the SWVR installation menu.
14. On Node A, run Cluster Administrator to confirm that the states of the cluster resources were restored to the desired point in time.
15. Start the cluster service on Node B. To do so, at Node B, run the `net start clussvc` command from a command prompt, or select **Computer Management**, and then **Services**.
16. From Cluster Administrator on Node B, verify that the cluster group can be moved between the nodes by right clicking the group, and then selecting **Move Group**.
17. Install SANworks Virtual Replicator:
    a) Insert the SANworks Virtual Replicator CD into the CD ROM drive.
    b) Select **Complete Setup**.
18. Reinstall the NetWorker client software on Node B.

Recovering from a Total Disaster

A site disaster such as a fire or flood could result in the loss of both nodes, both controllers and all or some of the data. This section assumes the following facts:

- Documentation of the controller configuration, as well as the configuration of SANworks Virtual Replicator resources is on hand.
- A backup of the System State and Quorum Disk is present.
- A backup of all data is present.
- All destroyed or damaged parts have been replaced.
- No disk devices are recoverable. If some disk devices are recoverable, and their respective RAIDsets or mirrorsets are still intact, then see the section titled "Recovery from a Storage Enclosure Failure and Data Loss" earlier in this paper.

This section will cover the procedures to recover from a total disaster.

1. Perform a QuickRestore on both Node A and Node B by selecting **Full**.
2. Reconfigure both controllers:
   – Set the worldwide name on controller A:

     `Set this node_id=` (Enter world wide name and checksum,)
   – Set the controllers for multibus failover:

     `Set multibus_failover copy=this` (This will restart the controller.)
   – Set the prompts for both controllers:

     `Set this prompt="top>"`

     `Set other prompt="bottom>"`
   – Set the time for the controllers:

     `Set this time=dd-mmm-yyyy:hour:min:sec` (For example, 14-feb-2001:15:45:01)

– Set battery expiration date on the top controller:

  Run FRUTIL, press the **Y** key, and then press the **Enter** key.

  Note: FRUTIL must be performed for both controllers. It will be necessary to switch the serial cable to the bottom controller.

– Set port topology:

  TOP> set this port_1_topology = FABRIC

  TOP> set this port_2_topology = FABRIC

  TOP> set other port_1_topology = FABRIC

  TOP> set other port_2_topology = FABRIC

– Run configuration to detect the disk devices by typing the following command:

  RUN CONFIG

– Create the Quorum mirrorset and all RAIDsets by completing the following:

  For a mirrorset, use the command:

  Add mirrorset M0 diskxxxxx diskxxxxx (where xxxx equals the disk number)

  For a RAIDset, use the command:

  Add RAIDset R1 diskxxxxx diskxxxxx diskxxxxx (Where xxxx equals the disk number, use the diskxxxxx parameter for all disks that must be in the RAIDset)

– Initialize the RAIDsets and mirrorsets that have been created:

  Init m0 (Where m0 is the mirrorset or RAIDset name)

– Create a unit for all of the disksets created using the following command:

  Add unit d0 m0 (This will create a unit called d0 and assign the mirrorset m0 to it.)

– Disable the port security setting for all units:

  Set d0 disable=all (Perform this command for all units.)

– Configure all port 2 connections for a unit offset of 100:

  Type Show Connections (This will show all connections to the controllers.)

  Set (connection name) unit_offset=100 (For example, if !newcon00 is connected to port 2, then type: Set !newcon00 unit_offset=100)

  Note: It is easier to keep track of your connections if you set them up one at a time and rename them. Further information on this procedure can be found in the installation guide.

– Restart Node A.

– Enable port security for the mirrorset:

  Set m0 enable= (Enter the name for all connections here. They must be separated by a comma with no spaces.)

On Node A:

1. Configure all network connections and domain configurations for both nodes.
2. Configure the Quorum Disk.

   Note: Logical Volume Manager (LVM) is disabled on TaskSmart N-Series appliances by default because it may interfere with the operation of SANworks Virtual Replicator (SWVR). SWVR is the appropriate disk management tool to use when configuring volumes; however, LVM is sometimes needed. To avoid the possibility of serious data corruption, be sure to never run LVM and Snapshot Manager at the same time.

   a) Open the MMC through the TaskSmart Console.

   b) Open the **Core Operating System** folder.

   c) Open the **Local Computer Policy** folder.

   d) Open the **User Configuration** folder.

   e) Open the **Administrative Templates** folder.

   f) Open the **Windows Components** folder.

   g) Open the **Microsoft Management Console** folder.

   h) Select **Restricted/Permitted Snap-ins**.

   i) In the right pane, double-click **Disk Management**.

   j) In the **Policy** tab (which should be displayed as default), click **Enabled**.

   k) Select **OK** to enable Disk Management.
3. Make sure the computer is logged onto Node A and logged onto the domain.
4. Right click **My Computer**.
5. Select **Manage**, and then **Disk Management** to start the Logical Volume Manager.

   The first three logical disks visible are the three partitions on the TaskSmart appliance mirrored operating system drives. In addition, a fourth logical disk should be present with a small partition. This is the Quorum Disk mirrorset created earlier in the HSG80 controller CLI interface. If the Quorum mirrorset does not show up in LVM, the administrator must rescan the disks.
6. Write a disk signature to the Quorum logical disk. The signature must be basic, not physical, in type.
7. Create a primary partition on the Quorum logical disk using all of the available space on the drive.
8. Format the disk partition with NTFS and the default allocation size.
9. Name the new volume **Quorum**.
10. Assign drive letter D to the Quorum volume.
11. Close the Logical Disk Manager interface.
12. To disable the Logical Disk Manager, repeat the steps from the Enable Logical Volume Manager section, except click **Disable**.

IMPORTANT: Do not repeat these steps on the second TaskSmart server.

13. Verify that you can see the Quorum Disk on Node B. If you can not see it through LVM, then do a hardware rescan.
14. Install cluster services on Node A.
    a) Select **Start**, **Settings**, **Control Panel**, **Add\Remove Programs**, and then **Add\Remove Windows Components**.
    b) Select **Configure** for Cluster Services.
    c) Configure Node A as the first member of the cluster.
    d) Install SANworks Virtual Replicator.
    e) Insert the TaskSmart N-Series Cluster Supplemental Software CD, and then select the **Setup** icon.
    f) Select **Complete**. A restart is required.

On Node B:

1. Configure all network connections and domain configurations for both nodes.
2. Install cluster services on Node B.
    a) Select **Start**, **Settings**, **Control Panel**, **Add\Remove Programs**, and then **Add\Remove Windows Components**.
    b) Select **Configure** for Cluster Services.
    c) Configure Node B to join an existing cluster.
    d) Install SANworks Virtual Replicator
    e) Insert the TaskSmart N-Series Cluster Supplemental Software CD and select the **Setup** icon.
    f) Select **Complete**. A restart is required.

On the controller:

Enable port security for the remaining units.

Set D100 enable=Connection names (All connections that will be using a unit offset of 0 should be granted access to units D00 through D07. All connections that will be using a unit offset of 100 should be granted access to units D100 through D107).

On Node A:

1. Open SANworks Virtual Replicator and reconfigure your resources to their state prior to disaster.
2. Select **Start, Programs**, and then **COMPAQ SANworks Virtual Replicator**.
3. Recreate all pools and virtual disks.
4. When a pool is created in SWVR, an SE pool resource will automatically be created in Cluster Administrator. It will be necessary only to create the file share, network name and IP address resources for each group in which the pool(s) are members.
5. Open Cluster Administrator.

6. Ensure that both nodes are active members of the cluster.
7. Right click the cluster group, and then click **Move Group**.
8. The cluster group should move to Node B.
9. Recreate all resources to their prior state.

Note: The SE Pool resources have been automatically created from the creation of pools in SWVR.

It will be necessary to recreate the virtual servers and file shares to their previous state.

Note: A virtual server consists of both an IP address resource and a network name resource.

10. Restore data from backup.

Conclusion

A disaster can be devastating to any network infrastructure regardless of its origin. The documentation that has been previously presented is provided to help in the recovery of various types of disasters. It is crucial to always have current backups of all data, and to store those backups in a secure offsite location. The System State and Quorum Disk should always be included in any backup plan, and the backup must be preformed on a regular basis to ensure accuracy. Documentation is also critical to a recovery and must be considered as important as any system backup.

Further documentation on installation or configuration procedures can be found in the installation guide, as well as the administration guide for the TaskSmart N-Series Cluster.

Further documentation on supported backup software and hardware solutions can be found in the backup guide for the TaskSmart N-Series Cluster.

Further documentation on Legato Networker 6.0 can be found at the following link:

www.legato.com/products/protection/networker/networker6/

Appendix A

Downloading and Installing the ClusRest Tool

This downloaded tool includes the files required to run the tool, a documentation file, and a readme file with general information about the downloaded tools.

To download this tool, use the following link:

www.microsoft.com/WINDOWS2000/library/resources/reskit/tools/existing/clusrest-o.asp

To download and install a tool, perform the following steps:

1. Click **Download this tool** at the top of this page.
2. In the **File Download** dialog box, select **Save This Program To Disk**.
3. Select a location on your computer to save the file, and then click **Save**.
4. In Windows Explorer, go to the location where the downloaded file is saved, double click the file to start the installation process, and then follow the instructions.

The downloaded file is a self-extracting executable (.exe) file. Running the file installs the tool and documentation on the system.

Uninstalling Tools

To uninstall the downloaded tool, perform the following steps:

1. In Control Panel, double click **Add/Remove Programs**.
2. Select the tool, and then click **Remove**.